Herschel Observers' Manual

HERSCHEL-HSC-DOC-0876

Version 5.0.3

(Completely revised and extended post-Operations version)

2014 March 7

Herschel Observers' Manual

Table of Contents

Preface

The Herschel Space Observatory is an ESA cornerstone mission that was launched on 14 May 2009, alongside the Planck cosmic microwave background mission. Originally known as FIRST (Far InfraRed Submillimetre Telescope) its name was officially changed in the year 2000 in recognition of the 200th anniversary of the discovery of infrared radiation by William Herschel in 1800. Herschel covered the range from about 55 to 672 microns (530-5000GHz) -- a region that is effectively totally closed to ground-based astronomy -- using a suite of three state-of-the-art instruments called PACS, SPIRE and HIFI.

Herschel was an observatory mission: that is, its time was distributed among the community instead of being used for a large-scale survey. It was also a consumables-limited mission - its useful life depended on the lifetime of the helium in the dewar that was used to cool the instruments and was expected to be in the range from 3.5 to 4 years from launch. The temperature in the cryostat started to rise at the end of Operational Day (OD) 1446 of the mission and the end of mission was declared early on OD-1447, just over 2 weeks short of completing 4 years in space; you can find information on the events that constituted the end of mission on the HSC web page (http://herschel.esac.esa.int/). If you have access to the ESA intranet (ESA users only), you should also look at the contents of the "Herschel - How it All Ended" video.

As an observatory mission its success thus depended on the quality of the science that the community carried out with it and how effectively the helium in its dewar was converted into science. The "helium into science" ratio was the principal deciding factor in allocating time with the Herschel Space Observatory.

Many aspects of the Herschel Space Observatory have been revolutionary. It is, thanks to its innovative design, the largest dedicated infrared telescope ever to be launched into space by a considerable margin. For the astronomer this has converted into high sensitivity and a spatial resolution a factor of 6 better than any previous far-infrared telescope launched into space, making Herschel a pathfinder mission in the far-IR. In fact, over much of its wavelength range Herschel has been limited in sensitivity mainly by the confusion from the background of faint, unresolved sources. This makes Herschel a revolution for astronomy in a range of the far-IR that had hardly been exploited before its launch. Herschel observations will have a huge impact on astronomy and on our understanding of the universe for many years to come.

This manual describes the observatory aspects of Herschel: the spacecraft and its performance; the mission; the space environment in which the Herschel Space Observatory has been operating (very different from previous missions such as IRAS, ISO and the HST); and use of Herschel - from how an observing proposal was received and treated, through to final archiving of the data. The aim is to give an overview of Herschel to the user, describing everything that an observer or someone wanting to use Herschel data needs to know at a superficial level; where deeper knowledge is required afterwards, the observer should go to the specific documentation for each system or sub-system (e.g. the individual instrument manuals, the Data Processing user manual, etc.) The aim is that simply by reading this manual, or by using it for reference, someone who is planning to work with Herschel data has enough information to decide whether or not to proceed and to have a clear idea how to start.

When this manual was first written for the Guaranteed Time Key Programme Call back in November 2006, the launch of Herschel was still 30 months away and knowledge of how the spacecraft and instruments would behave in space was theoretical. Similarly, some important elements of the Science Ground Segment were still in development. At the time of this revision we are now well into post-Operations, six months after the end of helium and have characterised all aspects of Herschel's performance and operation as thoroughly as possible. As a result, this manual has undergone a further deep revision to reflect a mature post-Operations reality.

Similarly, when originally written, the emphasis in all Herschel documentation was very much on Uplink -- what was sent up to the satellite and how it would get to be sent up to the satellite -- rather than downlink -- the data received from the satellite and its processing. The contents of this manual still very much reflect an Uplink bias. For detailed information of the downlink aspects of the Herschel mission you should check the documentation for HIPE (the Herschel Interactive Processing Environment) at: http://herschel.esac.esa.int/hcss-doc-11.0/ (for HIPE 11), from early 2014, for

HIPE 12 at: http://herschel.esac.esa.int/hcss-doc-12.0/ and, from early 2015, for HIPE 13 at http://herschel.esac.esa.int/hcss-doc-13.0/.

This manual is not intended to be a complete guide and reference book to all the intricacies of the Herschel mission. It is intended to cover the basics of Herschel for people who are new to the mission, or who want to find out a little more about Herschel without going into exhaustive technical detail. As such, a significant part of this manual is dedicated to mission history. Those who need deeper information are directed to the various instrument and software manuals and, to the Explanatory Supplement that will be published during post-Operations, of which a first draft is expected to be available in late 2014.

Chapter 1. Mission phases

Herschel flight operations have been divided into a series of phases from the moment of launch. Broadly, these are check-out (in the first six months after launch), routine operations (from then until the end of helium) and post-operations (from end of helium to the end of all Herschel activity), each with their individual sub-phases. In theory, each mission phase should have an exact, identifiable start and end point but, in reality, the requirements of operations and the differing needs of the three instruments have made the different mission phases blend progressively into each other, with slow transitions and no clear start and end point. Similarly, the HIFI anomaly that occurred almost three months into the mission, when a chain of events, triggered by an SEU (see Section 4.2.1.3), led to a power surge that disabled the primary power chain, meant that HIFI was delayed by about 6 months with respect to PACS and SPIRE in entering routine operations. The anomaly obliged HIFI to go back for a time and re-conduct check-out activities that had already been completed months earlier, using the recovered systems on the back-up power chain.

Overall, despite the HIFI incident, Herschel operations have ran smoothly, largely due to the success of the long and intensive pre-flight test campaign process, first at ESTEC and later at Kourou. As a result, some activities could be advanced considerably over the anticipated pre-launch schedule. Despite the arduousness of the post-launch check-out and commissioning plan, remarkably few of the 298 planned check-out activities failed for such a complex mission, thus requiring much less re-planning of in-flight tests than might otherwise have been expected and, hence, an earlier start to the first science observations than was anticipated pre-launch.

Figure 1.1. Roll-out of the launcher for the Herschel-Planck mission on 13 May 2009 with the dark, threatening storm clouds visible behind.

The only other serious instrumental issue to occur during the entire mission was the loss of one of the two matrices of the PACS 160 micron array, on February 16th 2013, just two and a half months before the end of helium, although, as the second matrix continued to function perfectly, this only had a small impact on science, reducing the sky area covered at 160 microns and thus reducing senstitivity in maps in this band by about 30% but, thanks to careful re-design of observations, the effects of the failure on observations was largely mitigated for most observers.

1.1. Completed mission phases

1.1.1. Early mission history

Figure 1.2. Launch of the Herschel-Planck mission on an Ariane 5-ECA at 13:12:02UT on 14 May 2009. The Herschel-Planck mission logo is clearly visible at the top of the fairing. The fine weather and clear blue sky contrasts sharply with the conditions for roll-out (Figure 1.1), or those prevalent only a few hours before launch as the VIPs were taken to the launch viewing area in heavy rain during an intense storm.

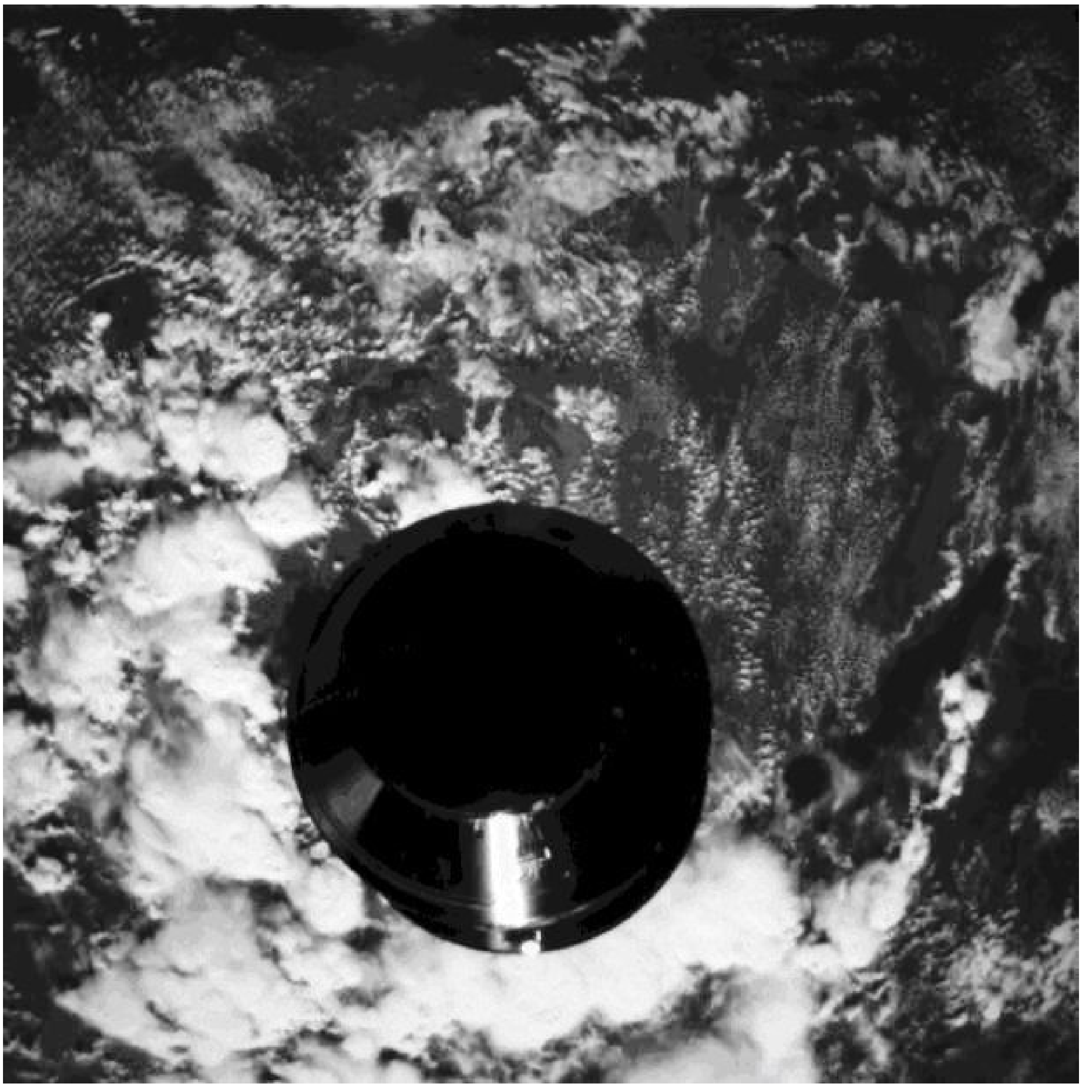

Figure 1.3. Herschel flies free! The moment of separation of Herschel from the Sylda, which is still attached to the upper stage and covering Planck. This image was taken with the backward-facing camera on Herschel: called the VMC. The coastline and cloudscape below Herschel are clearly visible.

Roll-out (Figure 1.1), prior to launch, was conducted early on the morning of 13 May 2009 and, after a flawless countdown, launch occurred at 13:12:02UT on 14 May, at the first possible opportunity as the launch window opened. Although there had been storms and heavy rain as the guests were being transported to the VIP area prior to launch, the clouds disolved and launch conditions were perfect. Figure 1.2 shows the Ariane 5 blasting off with the Herschel and Planck on board. The critical early milestones of fairing release, Herschel separation (at 13:37:55UT, Figure 1.3) from the Sylda and signal acquisition (at 13:49UT) all passed successfully. The frst command sent from the ground was executed by Herschel 58 min after launch at 14:10UT.

Launch occurred as the launch window opened, only 40 hours after the fairing was closed. This considerably reduced boil-off on the launch pad and gave a helium temperature at launch of only 1.81K, at the lower end of expectations, maximising the cryogenic lifetime in space.

After separation from the Sylda (Figure 1.3), Herschel began its cruise to Lagrange. A single injection manoeuvre was made 26h after launch. This injection manoeuvre started at 15:16:25UT and lasted 22.5 minutes, giving a Delta-V of 9.0m/s. This injection manoeuvre placed Herschel in what was very close to its final orbit. After the injection, a touch-up was made, 75 hours after launch, giving a further Delta-V of 1.0m/s. Two Orbit Correction Manoeuvres were made during transfer: the first, of 73cm/s, approximately 4 weeks after launch, the second of 17cm/s approximately 5 weeks after launch. No further corrections were required apart from the tiny, regular station-keeping burns of typically 20-30cm/s with the thrusters, made approximately every 8 weeks to maintain the orbit around L2 -- 21 of these station-keeping manoeuvres were executed in total during Herschel's 4 years around L2. Unlike for Planck, no braking manoeuvre was required to place Herschel in a stable manifold around L2.

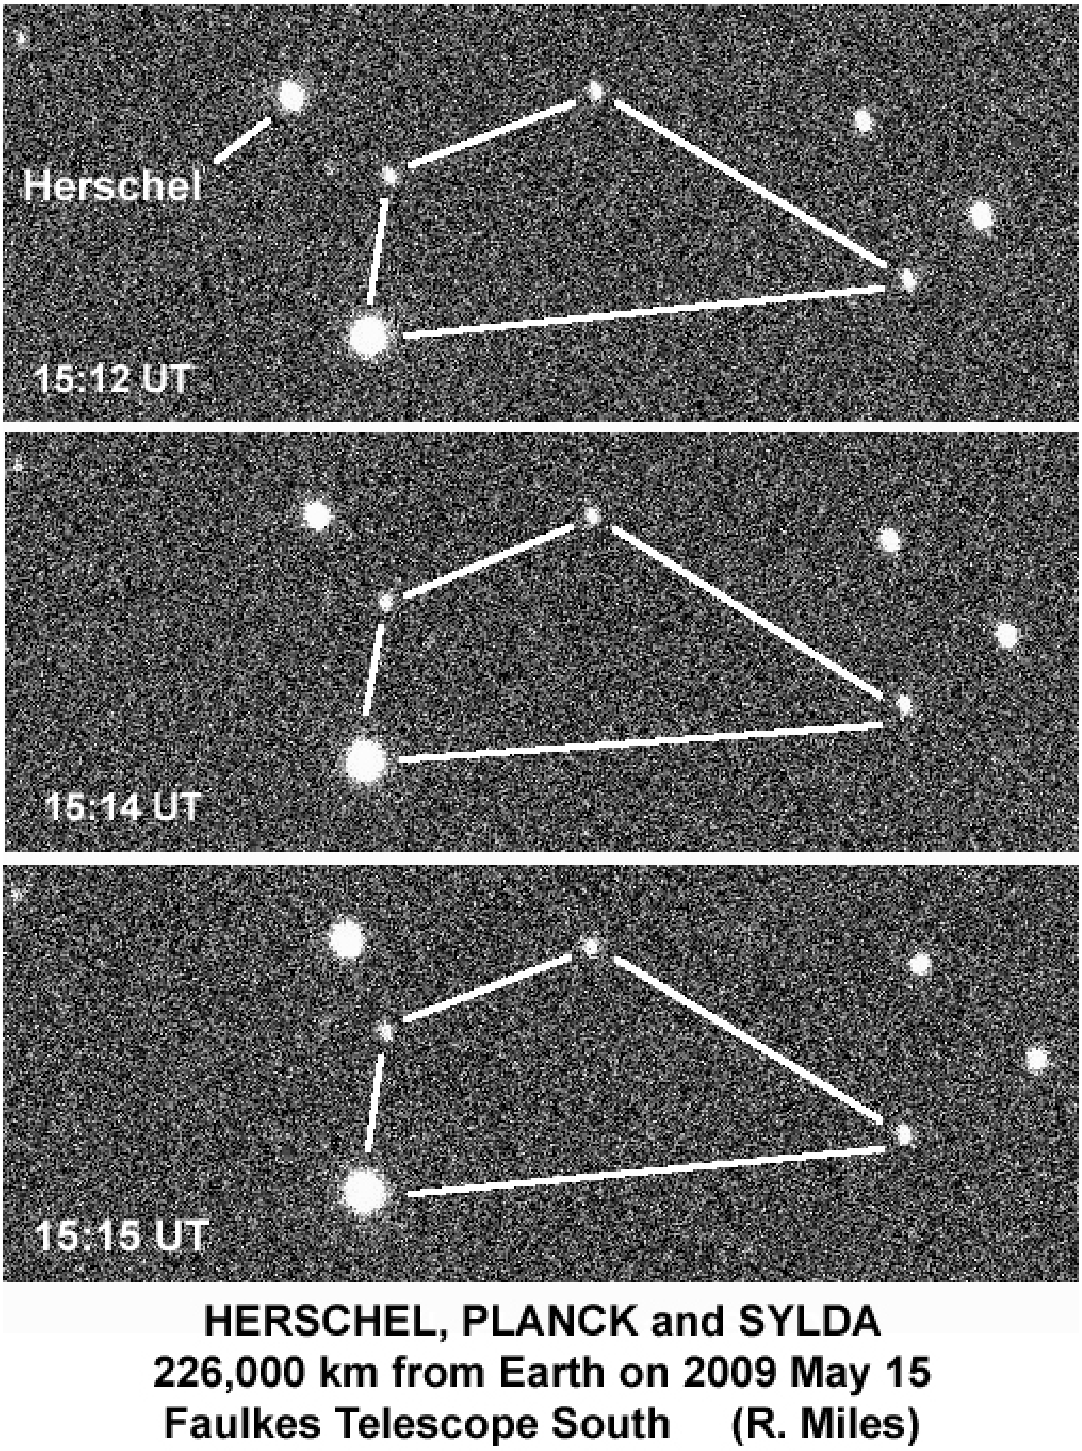

Figure 1.4. A sequence of images taken by British amateur astronomer Richard Miles using the 2-m Fawkes South Telescope in Australia, of Herschel (right), Planck (left, incorrectly labelled as Herschel) and the Sylda between the two, 26 hours after launch, at approximately half the distance to the Moon. These were the first observations of the Herschel-Planck constellation to be reported. The movement of the three relative to the stars in the three minutes between the first and last image is quite obvious.

Numerous images and animations were taken of the Herschel/Planck constellation as it receeded from Earth, mainly by amateur astronomers. Figure 1.4 shows the first image of Herschel, Planck and the Sylda to be reported from anywhere on Earth, taken approximately 26 hours after launch, when they were already 226 000km from Earth (approximately 0.5 Lunar Distances). This image was taken with the 2m Faulkes South Telescope. It is already obvious from this image how Herschel and Planck were released into slightly different transfer orbits due their differing injection requirements for L2 orbit. Excellent animations of the cruise phase of the Herschel/Planck constellation were taken by Peter Birtwhistle from Great Shefford Observatory in the United Kingdom (Minor Planet Center site code J95) on May 15th:

http://www.observadores-cometas.com/Herschel/images/Herschel_Planck_Sylda-090515-J95.avi

and by Gustavo Muler from Nazarth Observatory, Lanzarote, Spain (Minor Planet Center site code J47), on May 16th (from top to bottom in the animation: Herschel, Sylda and Planck)

http://www.observadores-cometas.com/Herschel/images/hershelplank_avi.gif.

The rapid variability of the tumbling Sylda is obvious, particularly in the Birtwhistle animation. It is also noticeable how the much larger Herschel actually appears considerably fainter than Planck in most of the images. Throughout operations in space Planck was consistently about a magnitude brighter than Herschel as it presented a greater effective surface area to Earth.

After launch, a Low Earth Orbit Phase (LEOP) started during the first stages of the cruise to Langrange. An initial phase of check-out started with the telescope closed and the instruments switched off, while the operation of the spacecraft sub-systems was checked. At this time the satellite was placed in a Sun-orientation, allowing the spacecraft to cool in the shadow of the sunshade and to outgas. However, given the danger of volatiles from the satellite condensing on the telescope, the mirror itself was heated during the initial cooldown phase to avoid it acting as a cold-trap for the outgassed volatiles. This period involved checking basic properties of the satellite (centre of mass, moments of inertia) and proper functioning of basic spacecraft sub-systems (Radio Frequency (RF), thermal control, power sub-system, data handling, attitude and orbit control, thrusters, Solid State Recorders (SSR), etc.) for spacecraft operations.

Figure 1.5. The telemetry received at MOC showing the oscillation in gyro response measured by the current to the gyros as the heavy cryocover swung open and oscillated, causing the entire satellite to wobble slightly until it had reached a stable open position. The large oscillation seen in this gyro signal (marked by the mauve coloured vertical lines on the monitor screen), showing the gyros activating to stabilise the satellite orientation, was the first evidence that the cryocover opening had been carried out successfully. Five or six oscillations occurred before the spacecraft stabilised completely again.

While spacecraft check-out was underway, initial activities to check-out the instruments could commence. The first instrument to be switched on to start payload operations was SPIRE on Day 6 after launch, as the heat generated by its electronics also helped to keep the temperature of the service module within its nominal operating range. PACS and HIFI switch-on took place on Day 11. Initial switch-on simply consisted of checking that the measured voltages were in the expected range from the telemetry. This led into a long and extremely detailed set of tests and checks of the functionality of each instrument that was the Commissioning Phase.

1.1.2. Commissioning Phase

Once Herschel was successfully launched and injected into the transfer trajectory towards the operational orbit, the spacecraft and instrument commissioning phase started. This consisted of a series of 298 individual tests and activities to check-out all aspects of instrument and spacecraft functionality. A highlight of Commissioning was the opening of the telescope cryo-cover. This cover protected the cryostat from condensation of outgassed volatiles. The cryo-cover was opened at 10:53UT on Sunday 14 June (12:53 local time at Darmstadt). This involved arming and then firing explosive bolts to free the cover, after which a spring pulled it into an upright, totally open position after a series of oscillations during which a spring steadied the heavy cryocover into a fully open position. The oscillations of the cover caused the gyros to activate to stabilise the spacecraft pointing as the cover was heavy enough to make the whole satellite move slightly in reaction to the movements of its opening, before finally coming to rest in an open position (shown in Figure 1.5). This gyro signal, showing five or six oscillations, was the first confirmation that the cryocover had opened, although it was not until the first images were received the following day that final confirmation was available that the cryocover opening had been a complete success.

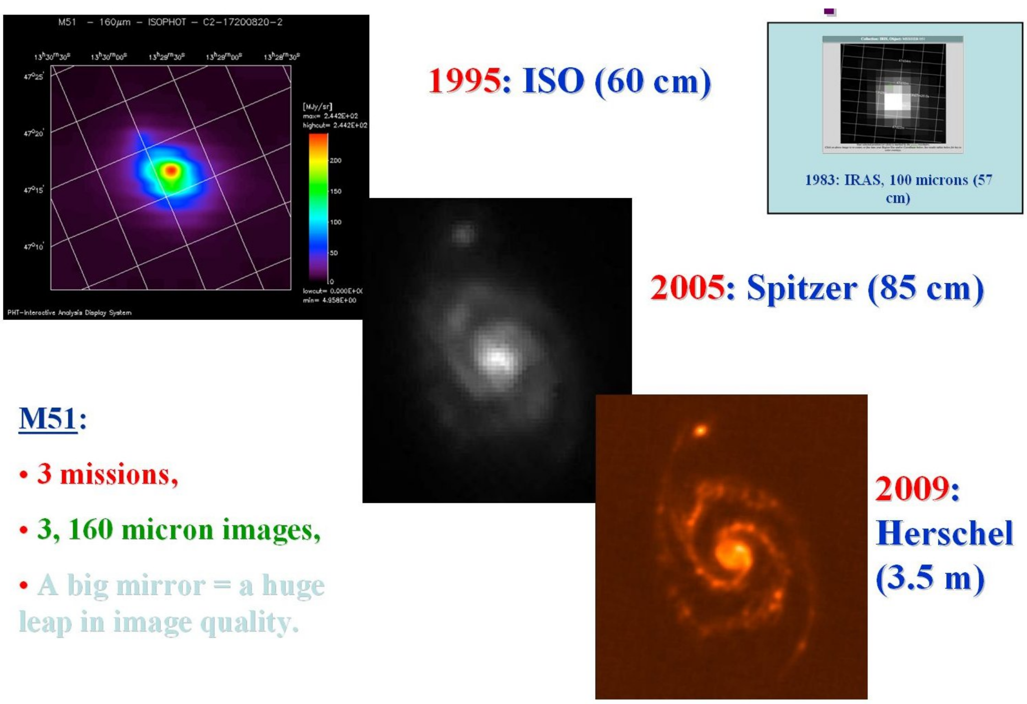

Figure 1.6. A comparison of images of M51 at 160 microns for ISO, Spitzer and the Herschel sneak preview image, showing the improved resolution and sensitivity from Herschel's larger mirror. No comparable image exists from IRAS, which had a long wavelength cut-off of 100 microns, so the IRAS 100 micron image is shown for comparison. The comparison of images also reflects the vast improvements in

detector technology over the years.

With the successful cryocover opening and the encouraging progress of commissioning activities, an opportunity was seen to take some early images to make a blind test of the telescope focus and image quality in advance of formal First Light. A series of PACS exposures were defined with a range of bias settings, scanning through the most likely range of values, in a test that was termed the "Sneak Preview". After consideration of several possible targets, the galaxy M51 was selected for the Sneak Preview for various reasons, especially the fact that it was available at the right time of year to be observable by Herschel and, being a large, bright object, the galaxy is a classic infrared target with a lot of structure, so considerable prior imaging existed at similar wavelengths with which the Herschel images could be compared (Figure 1.6) to check the image quality to ensure that it met expectations. The resulting images are shown in Figure 1.7, which showed that the telescope focus and alignment were excellent and that the optimum parameters for imaging were close to the best guess values estimated by PACS prior to cryocover opening.

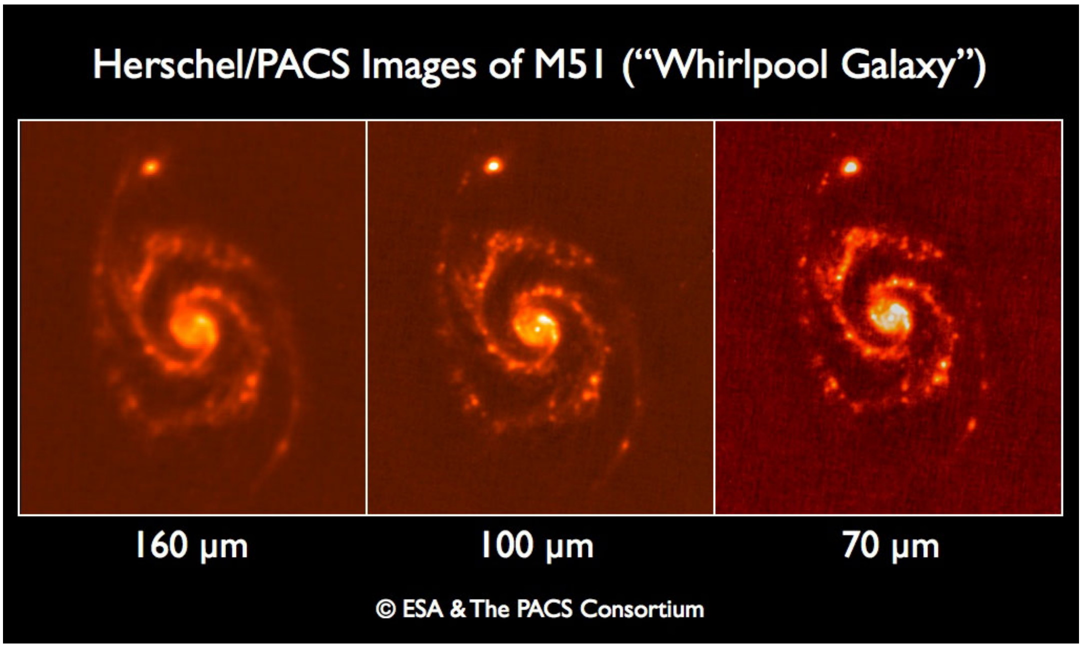

Figure 1.7. The Sneak Preview images of M51 in the three PACS bands, taken blind after cryocover opening on June 14/15th 2009.

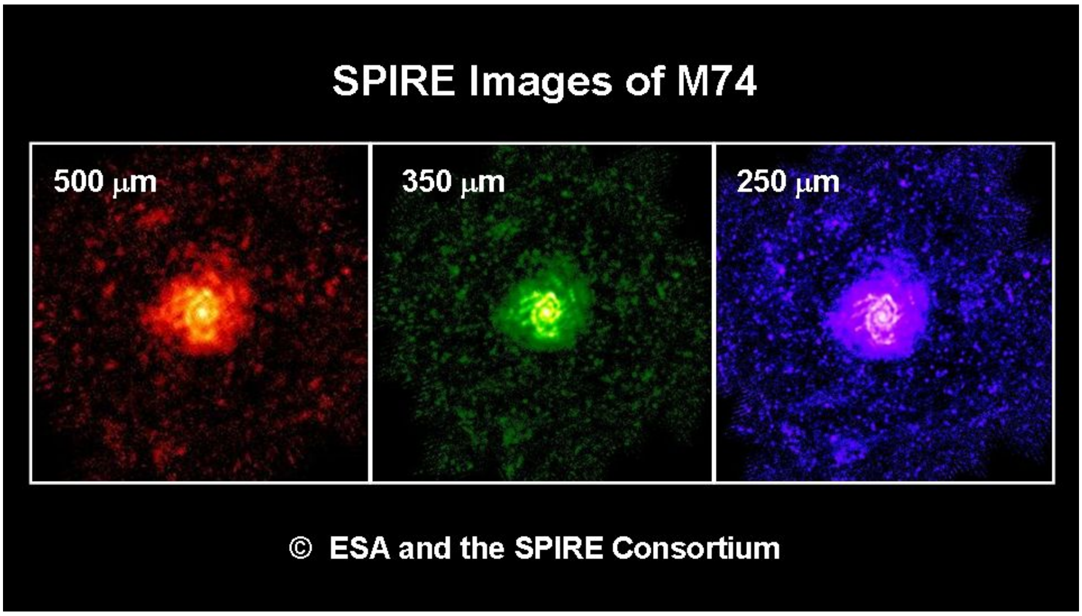

Figure 1.8. The SPIRE First Light images of M74 (NGC 628) in 250, 350 and 500 microns, obtained on 2009 June 24. Messier 74 is a face-on Sa galaxy in Pisces, at approximately 32 million light years distance with a visual diameter of approximately 10 arcminutes.

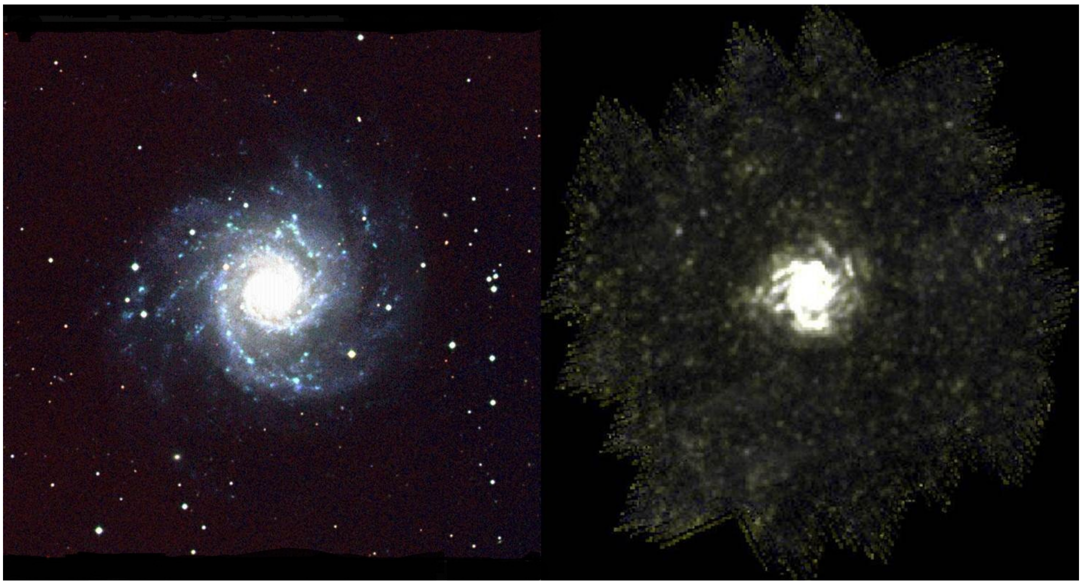

Figure 1.9. The SPIRE First Light image of M74 (NGC 628) combined as an RGB image (B=250 microns, V=350 microns, R=500 microns), obtained on 2009 June 24, compared with a Palomar Sky Survey RGB image on the same scale (blue layer = B, green layer = R, red layer = I) to show the comparison between the SPIRE first light image (right) and the Palomar Observatory Sky Survey image (left). This comparison shows where the star forming regions of the galaxy, bright to Herschel, are situated.

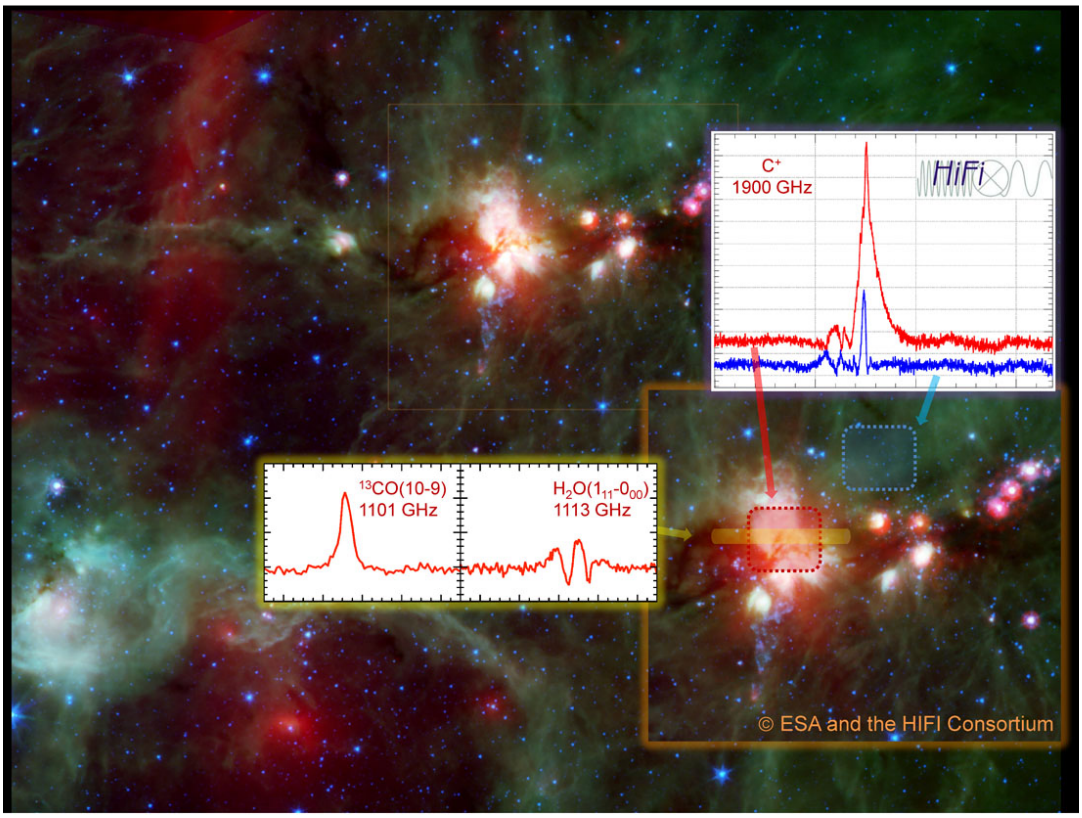

Figure 1.10. The HIFI First Light spectra of DR21, obtained on 2009 June 22, superimposed on a Spitzer image of DR21 and its surrounding region. In the inset we see an enlargement of DR21, which is part of a large star-forming complex in Cygnus, with the positions where the three HIFI spectra represented were taken, superimposed on the image.

Once the Sneak Preview had shown that the image quality was as good as had been anticipated in all senses, a more ambitious series of formal first light observations were scheduled for each of the instruments to demonstrate their capabilities, of which a sample is shown here (Figure 1.8, Figure 1.9, Figure 1.10).

1.1.3. Performance Verification (PV) Phase

PV phase was designed to obtain in-flight characterisation of all instruments e.g. in terms of stability, sensitivity, resolution, timing and other calibration parameters. It included the validation of the instrument observing modes and the calibration and data processing of the resulting data. To achieve this, a schedule of astronomical observations and internal calibrations, defined and iterated pre-launch, covering a nominal period of 2 months, were be executed using normal observatory procedures. This schedule was be based upon an agreed in-orbit calibration plan generated jointly by the ICCs and the HSC. The plan contained a description of all planned calibration activities and associated calibration sources (internal and astronomical) required to characterise fully each instrument.

Each instrument received blocks of time, normally of two days each, to carry out its activities according to the agreed PV plan, giving each instrument "two days on and four days off", allowing data to be processed and new observations prepared, albeit on an extremely intensive cycle. Weekly meetings then examined the progress of the planned tests, adjusting the plan to allow extra time for failed tests to be repeated, where necessary, or for extra tests to be included. PV Phase started 64 days after launch - in line with pre-launch plans - and, by 120 days after launch had delivered the first fully calibrated and usable observing modes for science scheduling.

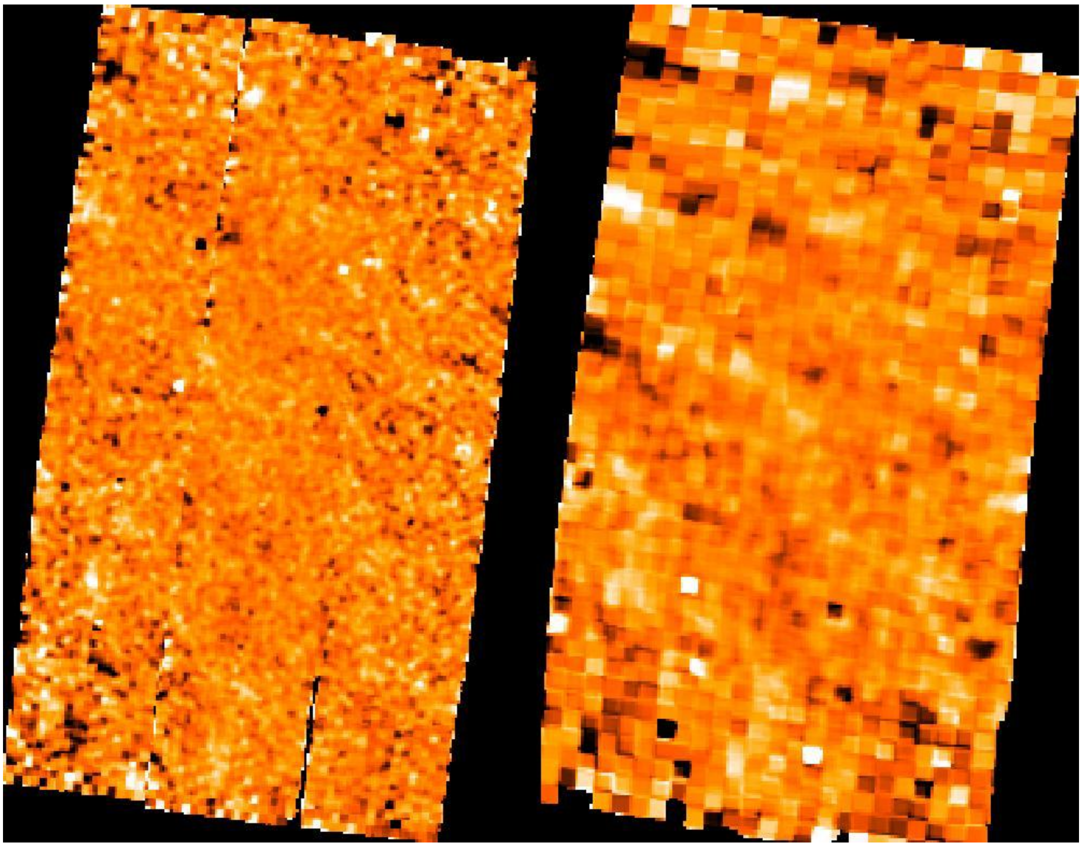

Figure 1.11. The browse product for the first successfully processed scheduled science observation made by Herschel. This is OBSID 1342183651 of the QSO SDSSJ1602+4228, made for AOTVAL_kmeisenh_2, a 2799 second exposure starting at 15:52:46UT on September 11th 2009 in Point Source Photometry mode showing an apparent faint detection at 160 microns (right hand image) where we see the weak pattern of negative and positive images from the chop/nod cycle that must be combined to give the final image and photometry. The left hand image was taken at 70 microns.

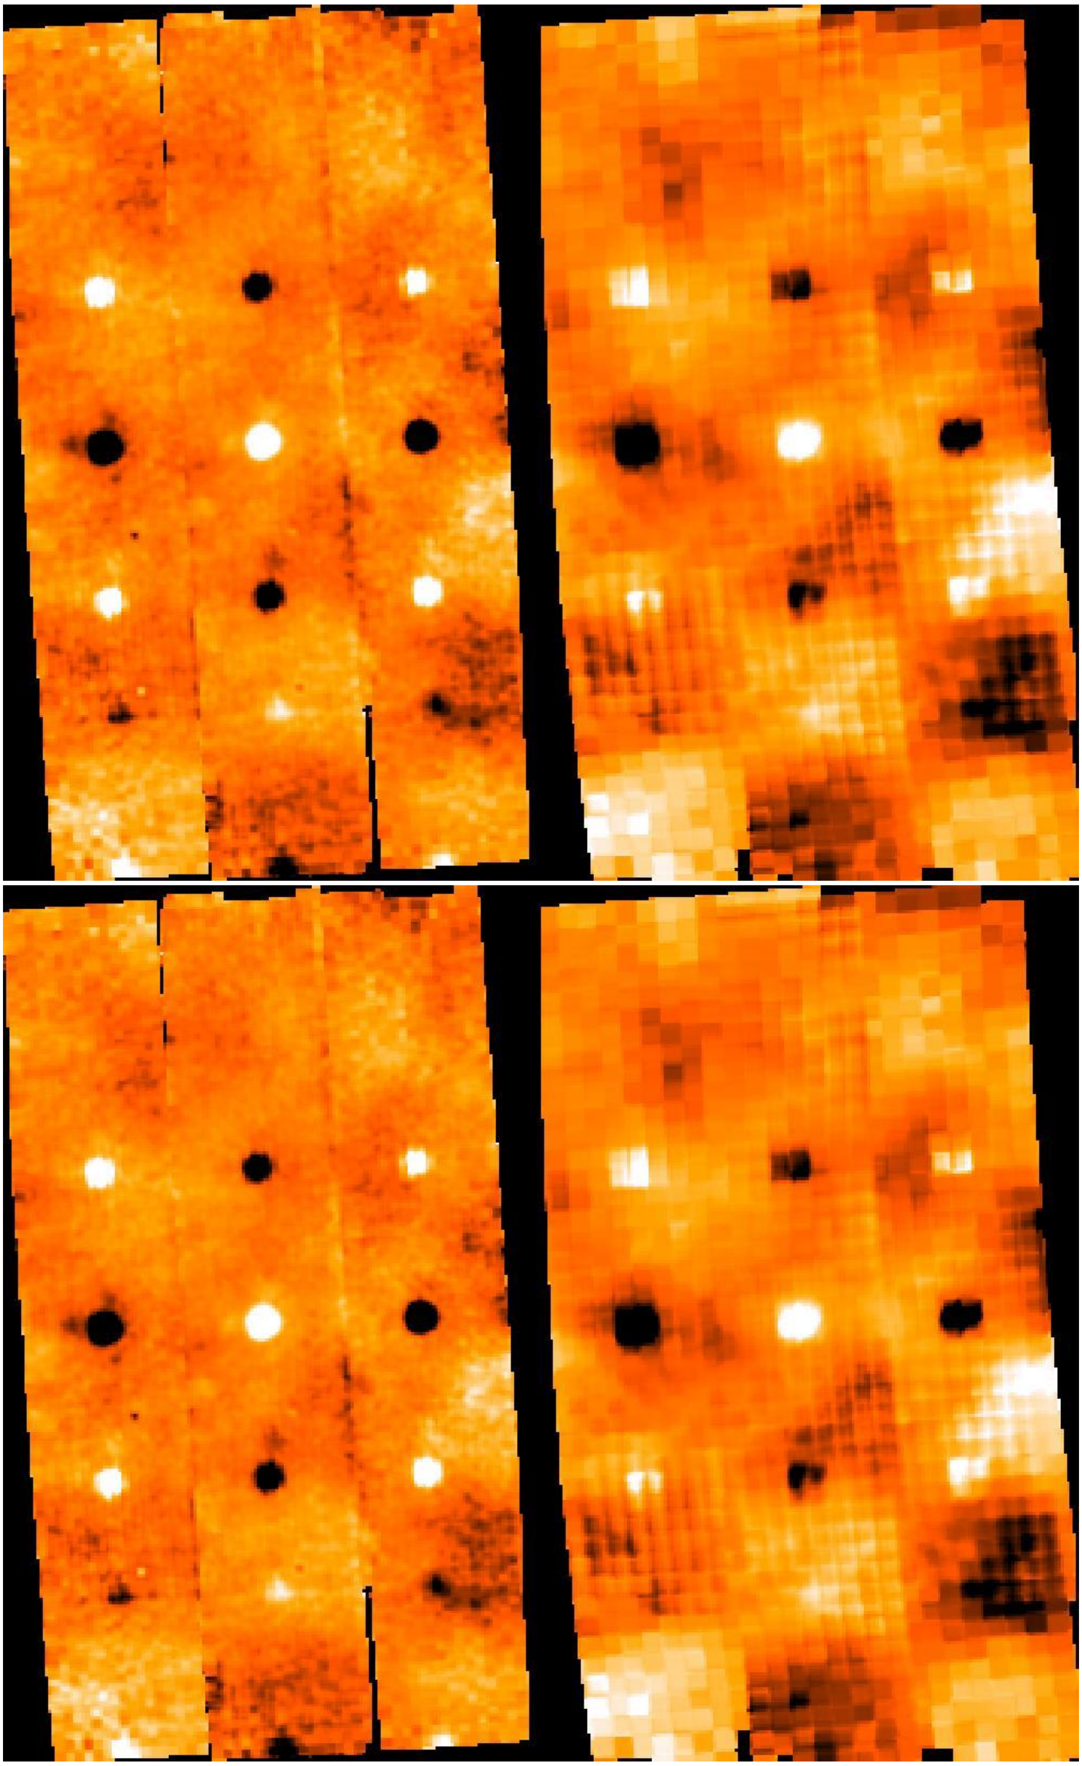

Figure 1.12. The browse product for the first scheduled science images taken of a Solar System Object

(SSO) made by Herschel. These are OBSIDs 1342183654 (top, 100 microns left, 160 microns right) and 134218365 (bottom, 70 microns left and 160 microns right) of Pluto, made for AOTVAL_thmuelle_2, starting at 17:01:46UT and 17:30:17UT on September 11th 2009 in Point Source Photometry mode. We see the pattern of negative and positive images from the chop/nod cycle on the individual photometer matrices, which must be combined to give the final image and photometry.

Once an observing mode was validated and considered ready for scientific use by the relevant Instrument Control Centre (ICC), a "Release Telecon" was held at which the observing mode was discussed in detail and all aspects of its readiness for release to the scientific community were discussed, including any final actions necessary to complete readiness for release, any documentation required and any caveats on use of the mode that should be communicated to the users. Once formal release was authorised, the Mission Planning Team at HSC would start to schedule routinely observations using that mode from the pool of approved observations in the database. In the early mission stages, to speed up scheduling of science observations, on some occasions an experimental version of an observing mode was scheduled to test and optimise it taking real science: these observations can be identified as they were labelled "AOTVAL", for AOT Validation and were taken on a shared risk basis, allowing some groups to obtain very early science data that could be used to refine their programmes.

PV Phase blended progressively into the Science Demonstration Phase of Routine Operations, without a formal end, with days assigned to PV activities becoming increasingly infrequent. Apart from HIFI recovery activities to check that the instrument was functioning correctly on its back-up chain after the incident in July 2009 that stopped observations, which ran in four dedicated blocks between 22 January 2010 and 17 March 2010, the last PV day included as such in the observing schedule was 8 December 2009. After this, remaining PV activities to test and validate remaining observing modes were absorbed into routine calibration activities for each instrument, as were the activities designed to test and validate the new, second-generation observing modes that were offered later in the mission.

A number of observations, taken as calibration observations, even relatively early in the mission during the PV Phase, have science quality and have subsequently been published. In some cases these calibration observations have been "bought" subsequently by a HOTAC-approved programme against its time allocation and thus promoted to science observations. However, the first science observation to be taken with Herschel that had been approved by HOTAC and was formally scheduled as a science observation, was a 159 second point source photometry observation of the protoplanetary disk J0843.3-7905 from the programme AOTVAL_bdent_2, cloned for the purpose from programme KPOT_bdent_2, starting at 15:28:21UT on September 11th 2009 (OD-120). This observation can be found in the Herschel Science Archive with OBSID 1342183650. A total of fourteen point source photometry observations were taken on OD-120, including AGNs (Figure 1.11), debris disks and a Solar System Object (Pluto, Figure 1.12).

One of the conclusions of the early observations in point source photometry mode was that the mode offered lower sensitivity than was hoped, leading to the mode being retired early for science use, although still used for calibration observations throughout the mission and, particularly, to carry out regular pointing calibration measurements to make a systematic check of the pointing quality. This mode was replaced by the mini-scanmap mode.

PACS scan maps were first used for scheduled science observations of HOTAC-approved programmes, in AOTVAL Mode in OD-127 (September 17th 2009), the first target being RCW 120 for AOTVAL_fmotte_2.

1.1.4. Science Demonstration Phase (SDP)

As PV Phase activities reduced, there was a progressive transition into "Science Demonstration Phase", in which each approved Herschel science programme had the opportunity to nominate a part of its observations -- typically 5-10% -- to be carried out early. The aim was to carry out observations and observing programmes that would test the capabilities of Herschel in detail, frequently with difficult and challenging observations. This allowed astronomers to test their observing strategy, compare data quality with expectations and fine-tune their observing programmes, as well as allowing a global overview to be obtained of the performance of the three instruments.

Observations were carried out on a shared-risk basis: astronomers could opt to forego their proprietary rights on data and allow them to be made public at the opening of the Herschel science archive and, in return, would get the time used re-imbursed in their programmes by Herschel; alternatively, they could maintain the data proprietary for one year from execution and have the data counted as part of their Routine Science programme. So, in return for assuming part of the risk of testing observing modes early in the mission, astronomers had the chance to obtain early publication of Herschel data, selecting their most critical observations for rapid execution.

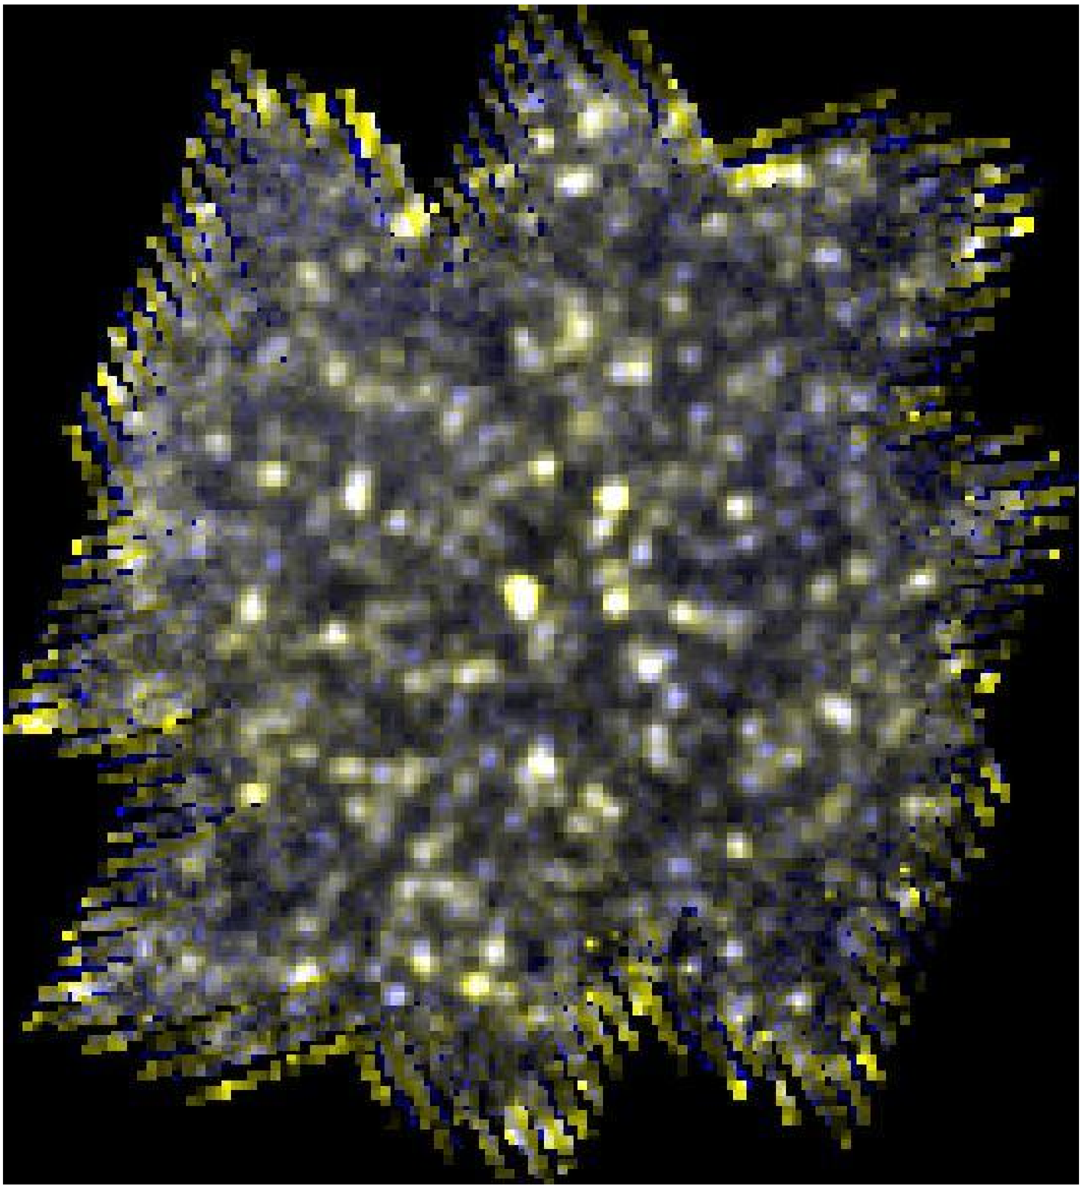

Figure 1.13. The browse product for the first SDP observation made by Herschel. This is OBSID 1342183677 of the QSO RXCJ0658.5-5556, made for SDP_eegami_3, a 6631 second exposure starting at 02:51:12UT on September 12th 2009 in SPIRE large map mode. This is an RGB image with the SPIRE 500 micron channel as the red layer, 350 microns as green and 250 microns as blue.

Science Demonstration Phase demonstrated that the main mapping modes that are the workhorses of Herschel, particularly in its early mission phases, were essentially ready to go, although some extremely useful input was obtained for observing strategies, leading to the recommendations on how to obtain the best sensitivity in observations. Similarly, it gave a lot of valuable information on how best to define spectroscopic observing modes. Early results from Science Demonstration Phase were presented at the Herschel "Initial Results" Workshop at Madrid in December 2009 at which every single Herschel observing project approved by HOTAC (42 Key Programmes) had some observational data to present.

The prime period for SDP was defined to be from 15 September to 15 December 2009. Those SDP observations that could not be completed by 15 December because of target visibility concerns, or because the required observing mode had not been released, were given the highest priority for telescope scheduling over the period up to 30 April 2010. After this date, the few remaining SDP observations that had not been scheduled for whatever reason, reverted to being treated as Routine Science observations.

As Commissioning and PV activities were advanced with respect to expectations, the first Science Demonstration Phase observations were executed even before the official start of the prime SDP period. OD-121 was entirely dedicated to SDP observations requiring SPIRE large maps. The first SDP observation to be executed was a SPIRE large scan map of RXCJ0658.5-5556 for SDP_eegami_3 (Figure 1.13), taking 6631 seconds to complete, which started at 02:51:12UT on September 12th 2009. Regular scheduling of SDP observations started on OD-145 (October 5th 2009), from which point blocks of SDP time were placed in the schedule with increasing frequency; from October 14th 2009 (OD-155) the schedule was dominated by SDP observations.

The first day dedicated to PACS Spectroscopy science observations was OD-165 (October 26th 2009). The first PACS spectroscopy observation was OBSID 1342186305, executed starting at 00:17:47UT for SDP_esturm_3, while the first SPIRE spectroscopy science observation was of Mark 231 for SDP_pvanderw_3, OBSID 1342187893 (Figure 1.14), starting at 07:17:14UT on December 9th 2009 (OD209).

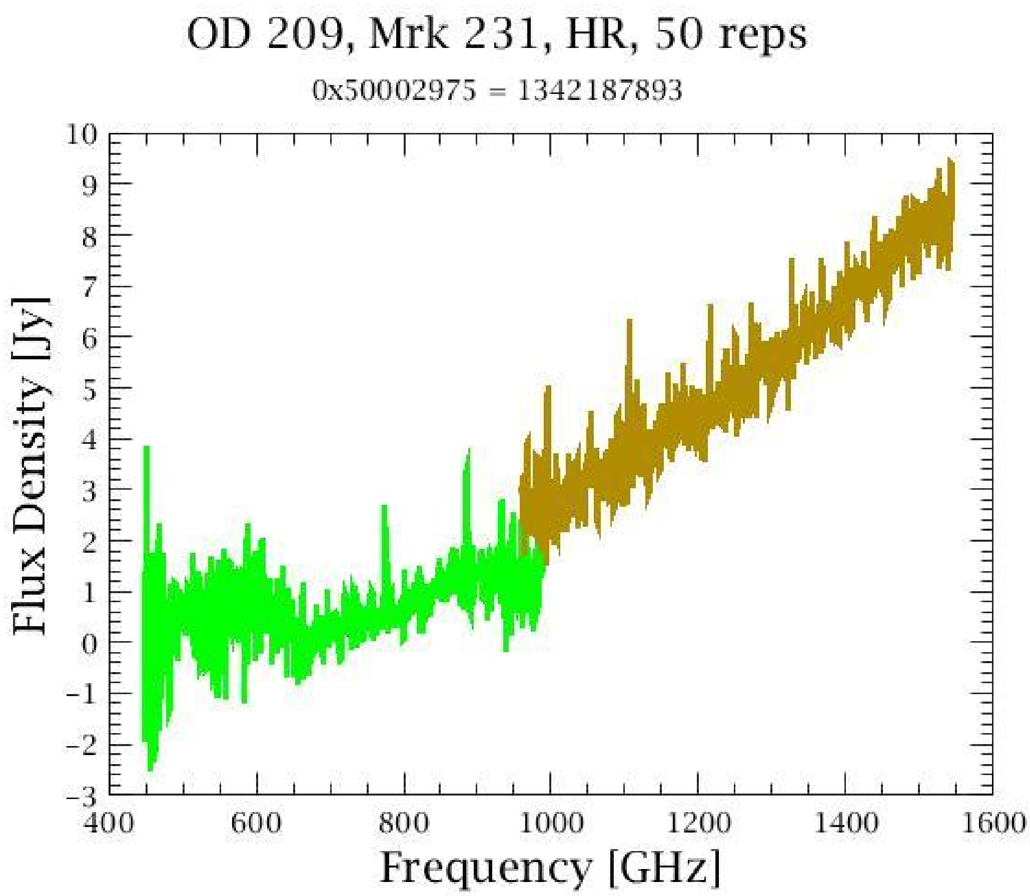

Figure 1.14. The browse product for the first SPIRE spectroscopy observation made by Herschel. This is OBSID 1342187893 of the AGN Mark 231, made for SDP_pvanderw_3, a 7141 second exposure starting at 07:17:14UT on December 9th 2009.

1.1.5. HIFI Priority Science Programme (PSP)

After the HIFI anomaly, it was decided to define a variant of SDP for HIFI observers, called the Priority Science Programme, to be executed as rapidly after HIFI recovery as time permitted. HIFI was switched back on on January 21st 2010 (OD-253) and placed in Dissipative II mode (lasers on), with this command being completed successfully at 18:31:31UT, allowing the lasers to stablise for two days of check-out activities on OD-254 and OD-255. Two blocks of five days of intensive re-commissioning and PV were carried out later in January and in February before HIFI science observations could start on February 28th 2010; the first HIFI science observation was a Band Ia observation of NGC 6334, starting at 22:16:18UT for SDP_cceccare_3.

All the initial HIFI observations were taken within the HIFI Priority Science Programme, or PSP. Two blocks of telescope time in March and April 2010 were reserved for intensive HIFI observing campaigns. Observations were divided into PSP1 (highest priority) and PSP2 (second priority) to fill these blocks of time efficiently, allowing a substantial part of the Herschel's HIFI observations to be carried out at the earliest possible date, after which HIFI would enter the standard Mission Planning cycle with a set number of days assigned each 4 weeks. Top priority, when HIFI re-entered the standard Mission Planning cycle was given to scheduling remaining PSP observations. A special HIFI initial results workshop was arranged in Leiden in April 2010 to present a first look at PSP data and checkpoint for observing strategies.

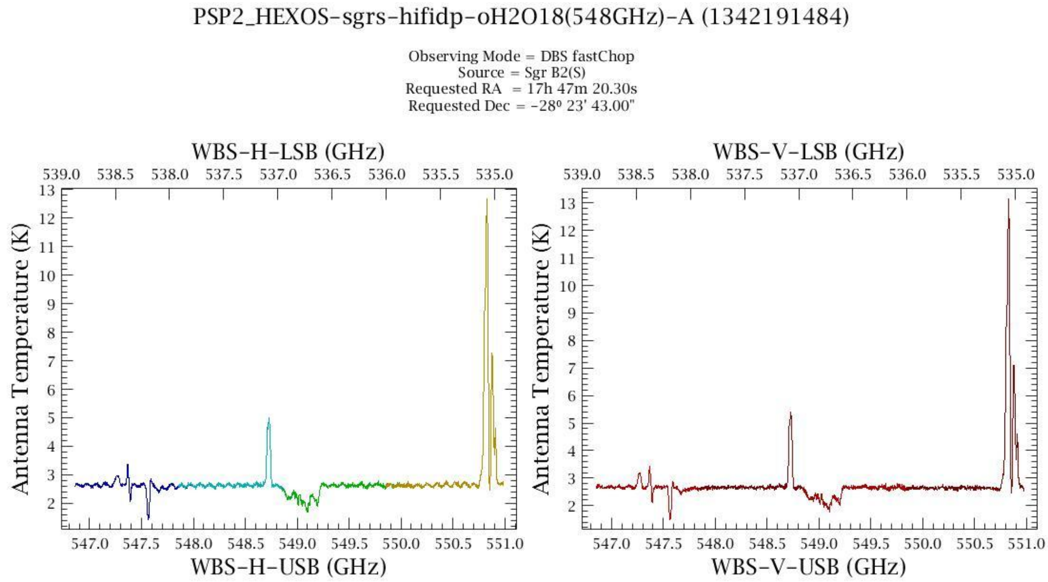

Figure 1.15. The browse product for one of the first HIFI science observations made by Herschel on the first day of the HIFI PSP programme (actually the 4th HIFI observation of PSP). This is OBSID 1342191484 of Sgr B2, made for KPGT_ebergin_1, a 61 second exposure starting at 06:29:20UT on March 1st 2010.

1.2. Routine Operations and Post-Operations mission phases

1.2.1. Routine operations (Routine Phase)

As with previous mission phases, there was a very gradual transition between SDP and Routine Operations. As each project received its SDP data, if no significant problems were revealed, the Principal Investigator (PI) was invited to have a release telecon with the Project Scientist and HSC staff to discuss the data and any problems or issues that had arisen. If no serious issues were identified, the PI was invited to release all, or part of the observations in his or her programme for scheduling, in which case, the observations would be made available to the HSC Mission Planners to be scheduled. The first routine observations were executed on 18 October 2009 -- a PACS scan map of the Lockman Hole for KPGT_dlutz_1. By December 2009, the immense majority of scheduled observations came from released routine programmes.

Over the course of the mission Herschel has produced hundreds, or thousands of spectacular images like the ones shown in Figure 1.16 and Figure 1.17.

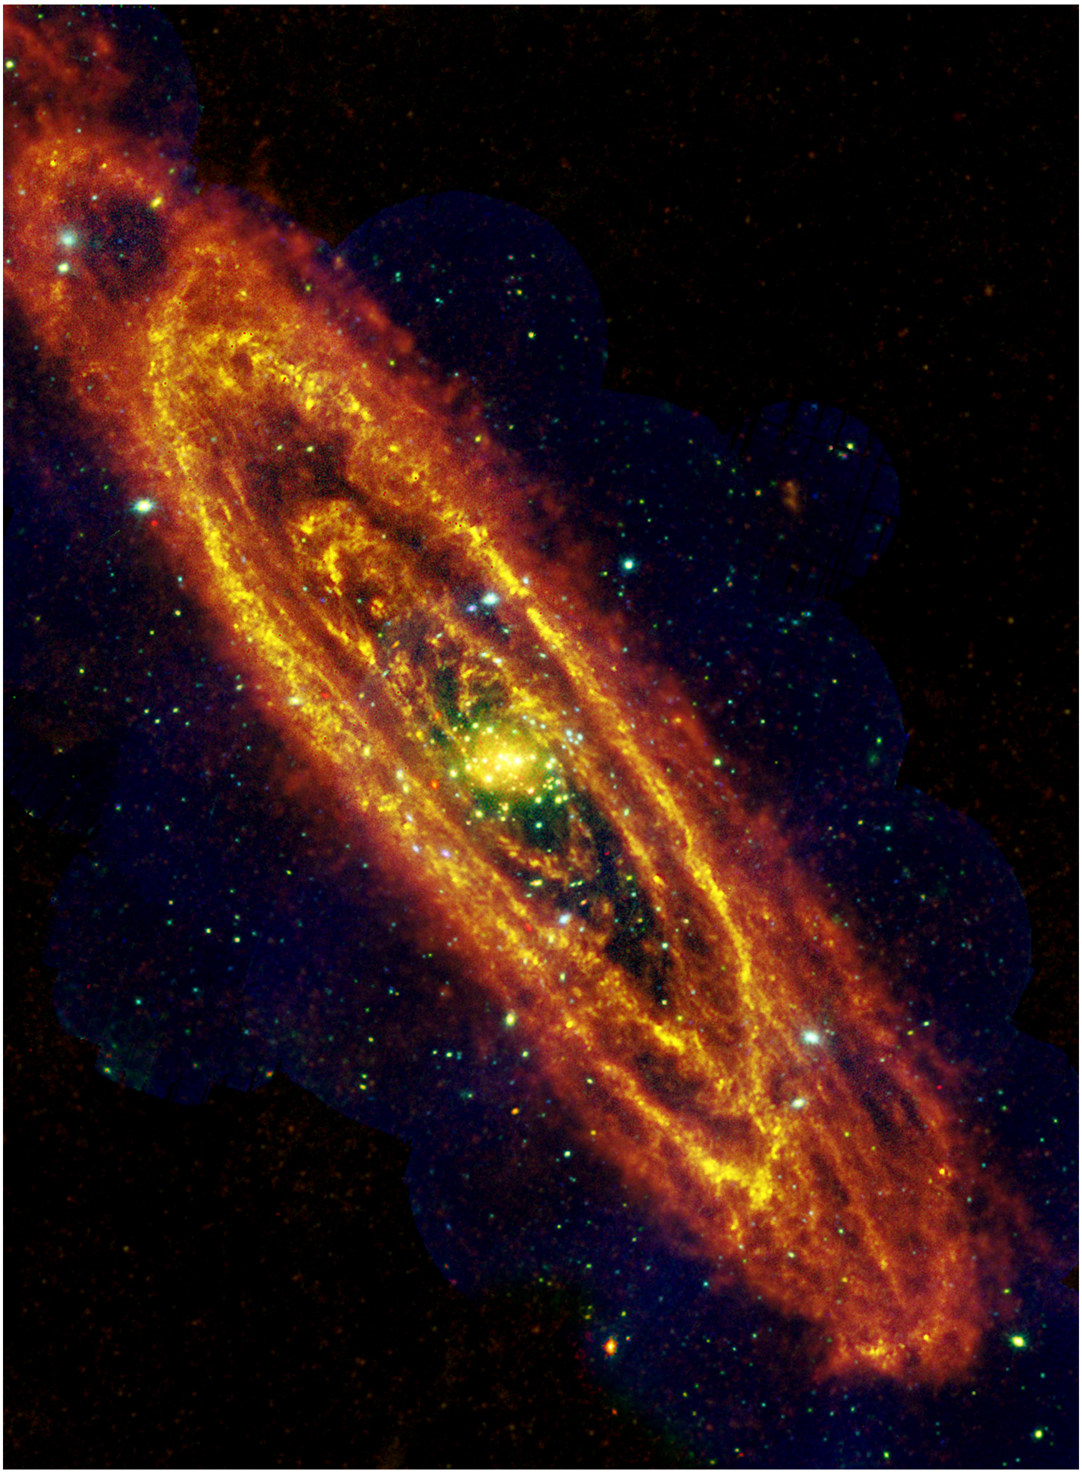

Figure 1.16. M31's once and future stars. A combined Herschel and XMM-Newton (x-ray) image of M31 showing dusty star-forming regions (Herschel) and the point-sources that represent highly evolved stars (XMM-Newton). The Herschel data were taken at 250 microns with SPIRE between 17 and 21 December 2010. In the XMM-Newton RGB image, red sources are those that have soft x-ray spectra, dominated by low-energy x-ray emission; these are normally low-mass x-ray binaries, while the blue sources are

sources with hard x-ray spectra, dominated by high-energy emission, which are high-mass compact binaries with a neutron star, or black hole secondary. This image contains what is effectively a snapshot of the star formation history of M31 and its two satellite galaxies, M32 (superimposed on the spiral arm below the nucleus of M31) and M110 (very faintly visible in the Herschel image to the top right). Whereas M110 is essentially a dying galaxy, with only a tiny residual of star formation and no massive stars at all, M31 is very much alive and boiling with star formation activity in the dark rifts between its spiral arms.

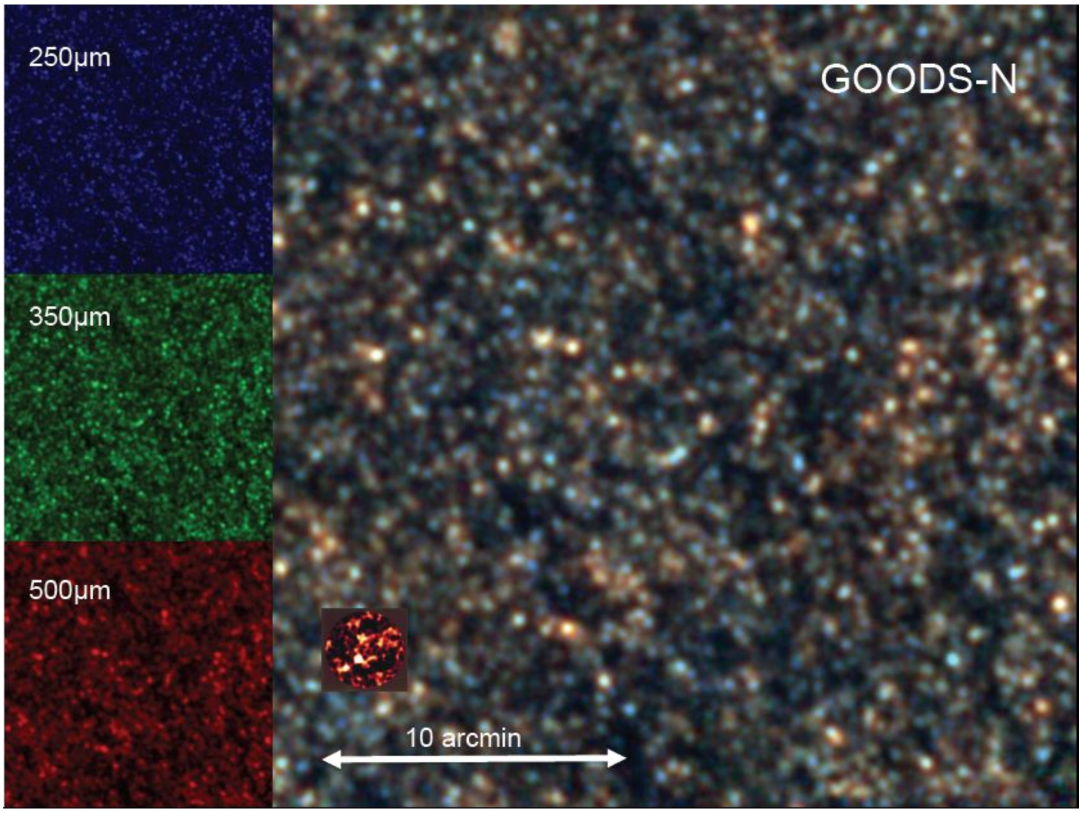

Figure 1.17. Galaxies spread like grains of sand on a beach. Every source in this GOODS-N field, which is about the size of the Full Moon, is a distant galaxy. The insert to the left shows the indivdual frames in each of the SPIRE bands while the main image combines them as an RGB. The colour of the galaxy gives an indication of its red shift and, hence, distance: the reddest galaxies are the most distant and may be as much as 12 000 million light years away; blue objects are relatively nearby and may be as close as 6000 million light years. In the inset to the lower left we see, to the same scale, the very best sub-millimetre image of this field that had ever been obtained from the ground at this time. This inset image is the result of 20 nights of exposure from Mauna Kea with the James Clerk Maxwell Telescope + SCUBA, in exceptional observing conditions; this image is to the same scale as the Herschel image - the full SCUBA field is 2.5 arcminutes diameter - and shows exactly 5 sources compared to Herschel's 15000 sources in the full 4x4 degree field of this region, detected in just 16 hours of exposure.

The pre-launch goal was for Herschel to carry out routine science operations for a minimum of 3 years, with an aim to complete 20 000 hours of science observations. To maximise the science return this time was carefully divided up into segments. During the first two years of the mission, Guaranteed Time and "Key Project" observing programmes received absolute priority. Key Projects were performed early in the mission to permit follow-up and to give the Guaranteed Time holders the opportunity to obtain real data to work with, in preparation for supplying community support to the open time observers with the benefit of a thorough knowledge of the entire observing chain from proposal submission to access and reduction of data. Almost all Key Programme observations were completed by May 2011, at which point OT1 observing programmes started to be heavily scheduled, although some OT1 observations had been scheduled as early as December 2010 where they helped to improve observing efficiency by filling inconvenient gaps in the telescope schedule. A similar pattern, shifted by one year, applied to OT2 observations. All Priority 1 observations in the database, save for cases of force majeur and a few time-constrained observations were completed by the end of 2012. Over the last six months of the mission, the schedule was increasingly dominated by Priority 2 observations, as the pool of available Priority 1 was depleted.

All observers could track the state of their proposals from the (password protected) proposal handling pages of the HSC Web page and were notified when the resulting data had been passed through the initial Quality Control process; this would take a minimum of 2-3 weeks to complete as it involved an individual inspection by eye of all images and spectra by an expert at HSC. Even so, data was available for retrieval from the HSC usually within 48 hours of the observations being executed.

Observers could also check both what observations were scheduled for observation and which had been delivered to MOC, but not yet executed:

http://herschel.esac.esa.int/observing/ScheduleReport.html

and the log of executed observations:

http://herschel.esac.esa.int/observing/LogReport.html

from the HSC, which constitutes the formal Observing Log of executed science observations (engineering and calibration observation were not recorded in either the observing schedule, or in the Observing Log).

Observations declared to be "Failed" due an anomaly in execution were automatically cloned manually by the Community Support Team and released for re-scheduling by HSC, without the need for intervention by the affected astronomer, who was informed as soon as the failed observations were cloned and re-released for scheduling. Frequently, this process was completed before the astronomer was even aware that there had been a problem with the original execution of the observations.

1.2.2. Boil off

There was considerable uncertainty as to when boil off would happen. Much of this uncertainty was due to a lack of detailed knowledge of exactly how much helium Herschel carried at launch, because the boil-off after fairing closure was essentially unquantifiable. An attempt to estimate the lifetime was made by taking three Direct Liquid Content Measurements (DLCMs) of the cryostat in space. These used a small heat pulse into the cryostat to measure the temperature rise and, hence the mass of helium with a nominal uncertainty of approximately 5%. However, there was a significant disagreement between the three derived estimates leading to considerable uncertainty in the lifetime that was obtained by extrapolating the least squares fit between the three measures to zero (knowledge acquired a posteriori suggests that the second DLCM gave a value that was rather too low).

This uncertainty was significantly reduced by using data from cooler re-cyclings, with the temperature rise as the cooler was cycled giving a direct estimate of the mass of remaining helium (effectively, each cooler re-cycling was a mini-DLCM). The best estimate in late 2012 placed the most likely date for helium exhaustion around mid-February 2013, but with an uncertainty of several weeks. Further estimates, with data at increasingly low fill levels, extended the likely lifetime to late March and then, finally, into early April, at which point it became obvious that there was a small, systematic offset in the values derived by the models, although the fact that the temperature rise during re-cyclings was much larger than even six months previously and increasing rapidly with each successive re-cycling, made it evident that helium exhaustion was very close.

There was some concern that, in the last few days or weeks of the mission, the behaviour of the instrument coolers would change after the helium level had dropped below a certain point and that, possibly, the cooler hold time would suffer a significant reduction. This was uncertain at the time, as we did not know exactly what the helium behaviour would be at very low levels. However, the cooler behaviour was found to be close to nominal until the end of the cold mission. The penultimate cooler re-cycling, which was executed on OD-1440 (April 22nd), was somewhat anomalous in that the temperature during the re-cycling did not reach the expected level but, in the end, the hold time was only fractionally reduced from normal. In contrast, the final cooler re-cycle to be successfully executed during the mission, which was carried out on OD-1443 (April 25th), was completely nominal in all respects (helium boil-off occurred around two hours before the next scheduled cooler re-cycle -- a SPIRE re-cycle -- was due to start).

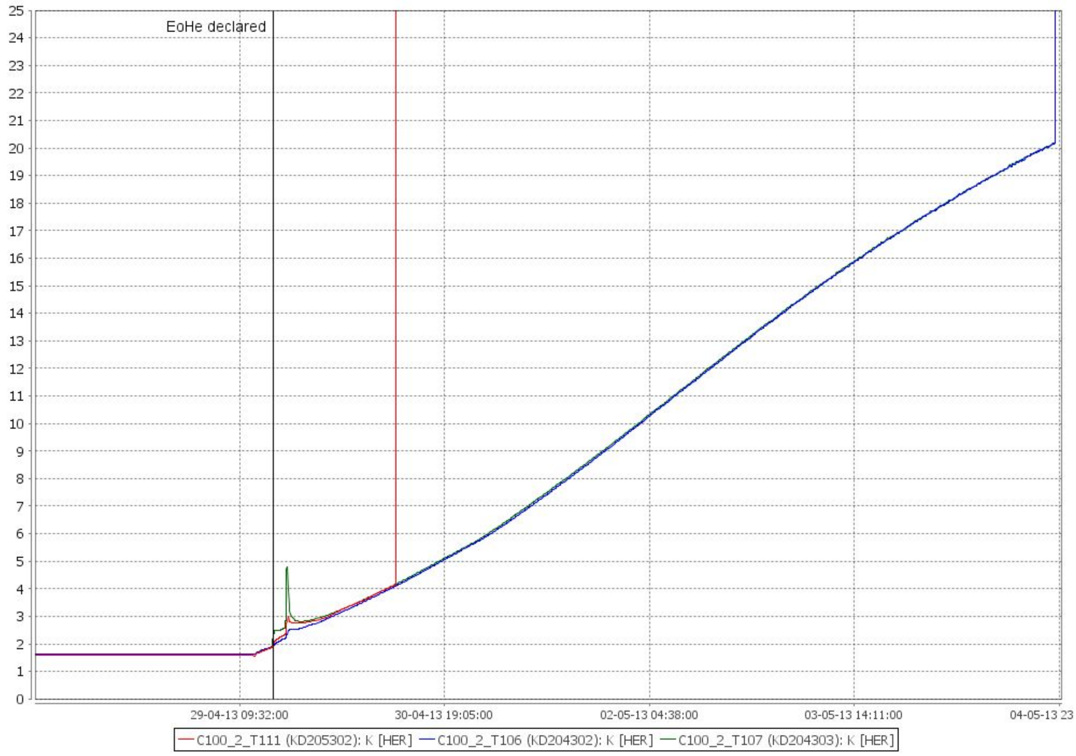

Figure 1.18. Cryostat temperatures from April 28th to May 4th 2013. The graph is initially stable at 1.7K before starting to rise around 12:30UT on April 29th, marking the evaporation of the last drop of helium. The solid black line marks the point at which an internal temperature of 2.0K was reached, marking the formal criterion for declaring End of helium. The plot shows the slow rise of the cryostat temperature after End of Helium (EoH) was declared. Temperature curves are shown for the T111 (red curve), T106 (navy blue curve) and T107 (green curve) sensors, which were in different positions inside the cryostat. As each sensor went out of limits it shoots vertically off the scale. Approximately five days after boil-off the temperature reached 20K and the T106 and T107 sensors went out of limits simultaneously (Graphic courtesy: MOC).

At Acquisition of Signal (14:15UT) on April 29th 2013 (the start of OD-1447) the temperature inside the cryostat was noticed to be anomalously high and trending upwards. This trend had started at 12:30UT approximately and marked the point at which the last drop of helium evaporated. Formal End of Helium (EoH) was declared at 15:20:01UT when the temperature within the cryostat passed 2K. Figure 1.18 shows the behaviour of the cryostat temperature as measured by three of the internal sensors, with the solid vertical black line marking the point at which the end of helium criterion was met; the other vertical lines mark the point at which each sensor went out of limits and subsequently registered as off scale.

The final, successfully executed science observation of the mission was completed at 07:40:57UT on April 29th. This was a HIFI Band 4b observation of the star OH 32.8-0.3 taken for DDT_kjusttan_3 and shows a strong detection of the OH line (Figure 1.19). As there was a Spacecraft Operations (SOPS) window for a manoeuvre scheduled after this observation, followed by a cooler re-cycle, no science observations were affected by the rising temperatures before EoH was declared and the instrument sub-schedule was disabled, stopping the telescope from continuing to attempt to execute observations.

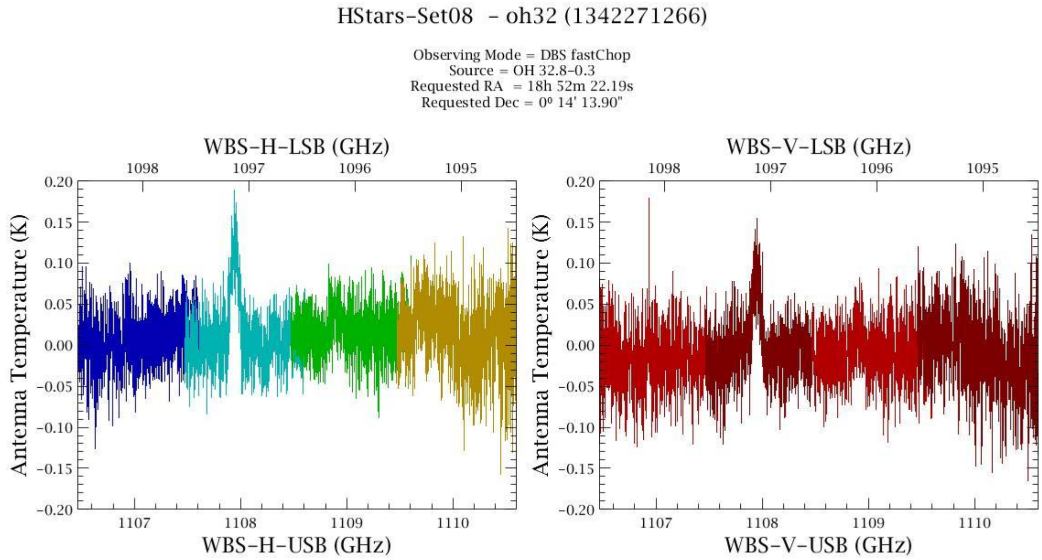

Figure 1.19. The final science observation executed before End of Helium. There is a strong detection of the OH line at 1107.9GHz (270.6 microns) in the HIFI Band 4b spectrum of the star OH 32.8-0.3. The two panels show the horizontal (left) and vertical (right) polarisation component of the spectrum.

1.2.3. Summary of mission phases and approximate dates

Table 1.1 gives a summary of the key dates in the Herschel mission. For the reasons explained above, it is extremely difficult to put hard dates for the beginning and end of some mission phases as the tended to have a smooth transition, for example, as in the case of the transition from Performance Verification to Science Demonstration to Routine Operations.

Table 1.1. A summary of Herschel mission key dates. Only approximate dates can be assigned to many of the different mission phases as there is inevitably a progressive transition between mission phases rather than a sharp one; on some occasions there were activities from three different mission phases progressing simultaneously and, in some cases, the start and end of a phase is a matter of definition and different dates could be given to the ones that appear here. In particular, HIFI recovery activities meant that CoP and PV days were scheduled months after the nominal end of these phases. Similarly, as reflected by this table, occasional PV days were being scheduled for PACS and SPIRE long after even routine observations had started.

Mission phase	Approximate Start	Approximate End
Launch	L=14 May 2009	
Early Orbit Phase	L	24 May 2009 (L+10 days)
Commissioning Phase	L	July 19th (L+66 days)
Performance Verification Phase	17 July 2009 (L+64 days)	25 November 2009 (L+195 days)
Science Demonstration Phase	11 September 2009 (L+120 days)	30 April 2010 (L+352 days)
Herschel Routine Phase	18 October 2009 (L+157 days)	L+47.5 months; Boil-off = B
Run-down phase (3 months)	B (April 29th 2013)	B+3 months (end July 2013)
Mission consolidation phase (6 months)	B+3 months (August 2013)	B+9 months (February 2014)
Active archive phase (48 months)	B+9 months (February 2014)	B+51 months (July 2017)
Archive consolidation phase (6 months)	B+51 months (July 2017)	B+56 months (End of Herschel mission), December 31st 2017
Historical archive phase (indefinite)	B+56 months (January 1st 2018)	(TBD) End of all Herschel activity

1.2.4. Post-Operations Phase

The Herschel post-operations phase consists of the rundown monitoring phase (starting at the moment of helium boil-off), mission consolidation phase, active archive phase, and the archive consolidation phase (at which point the transfer to the subsequent historical archive phase takes place), which is the final formal phase of the mission. Herschel is funded for almost 5 years of post-operations.

1.2.4.1. Observatory rundown, disposal and passivation

Once boil off had occurred, Herschel's instruments were no longer able to operate. As all the instruments needed temperatures below 10K to operate there was no possibility of a warm phase, as the temperature inside the cryostat was expected to warm within a few hours to 50K (in fact, the temperature took around four days to reach 20K). There was considerable discussion about what to do with Herschel finally post-helium. A key feature of the post-helium plan was to ensure safe disposal of the spacecraft

MOC continued to operate Herschel for seven weeks after boil off had occurred. This time was used for routine spacecraft housekeeping operations and a long series of technology tests involving all three instruments and various spacecraft sub-systems such as the star-tracker. Similarly, the radiation monitors continued to function as a space weather station until end of the mission (see Figure 4.7). A proposal that Herschel be crashed deliberately into the Moon was rejected and it was finally decided to place Herschel into a non-return solar orbit using a series of disposal manoeuvres.

The first stage of disposal was a short burn to leave the stable envelope around L2. This was done in advance of End of Helium. A burn of 10.51 m/s was made on March 15th 2013, at the end of OD-1401. This ensured that Herschel would start to drift away from the Earth (Figure 1.20)

The second stage of disposal was the main burn. This was made on May 13th 2013 (OD-1461), starting at 18:20:35UT and involved firing the thrusters for 7h46m to give a Delta-V of 113.7m/s. This was the longest burn ever made with thrusters of the type installed on Herschel (the EADS CHT20). This burn raised the aphelion of Herschel.

The length of the burn was determined by the need to ensure that it ended before the hydrazine in the thrusters was exhausted to ensure a controlled end of burn; this meant stopping the burn when it had consumed the smallest amount of hydrazine that it was believed could still be in the tanks.

The final stage of disposal was carried out on June 17th, at the start of OD-1497. With all tests complete, the satellite was prepared for final de-activation. The fuel was drained with a final manoeuvre to empty the tanks completely, a process that took approximately two hours to complete, in line with the best estimate of the amount of fuel remaining before the drain manoeuvre. This gave a final impulse of 31m/s. As the fuel was exhausted, the telemetry in the Operations Room at MOC showed that the pressure in the reaction chamber was dropping non-linearly (Figure 1.21). Once the fuel level became critical, the burn became unstable as different thrusters continued to fire weakly while others flamed out, at which point the burn was terminated manually and the spacecraft was stabilised with the reaction wheels, ready for the passivation. This consisted of sending a command to switch off the spacecraft transponder, so that it would no longer try to establish daily contact with Earth (Figure 1.22).

Once the command was sent, it took 6s to reach Herschel and a further 6s for the response of the transponder to reach Earth. It then took several seconds for the signal to drop to zero, indicating that Herschel was finally silent (Figure 1.23). With the transponder silent, the telemetry on the screens in the Main Control Room flatlined (Figure 1.24).

Herschel's final orbit

Initial assessment of the orbit after the final drain manoeuvre showed that Herschel was in an orbit completely exterior to the Earth's, with perihelion at 1.036AU and aphelion at 1.064AU. This will give a close encounter with Earth approximately every 14 years, with the next encounter in October 2027 at a distance of 0.050AU, equivalent to 7.5 million kilometres or 19.5 Lunar Distances, at which time the satellite will be magnitude 23 approximately and is likely to picked-up in routine monitoring for Near Earth Asteroids.

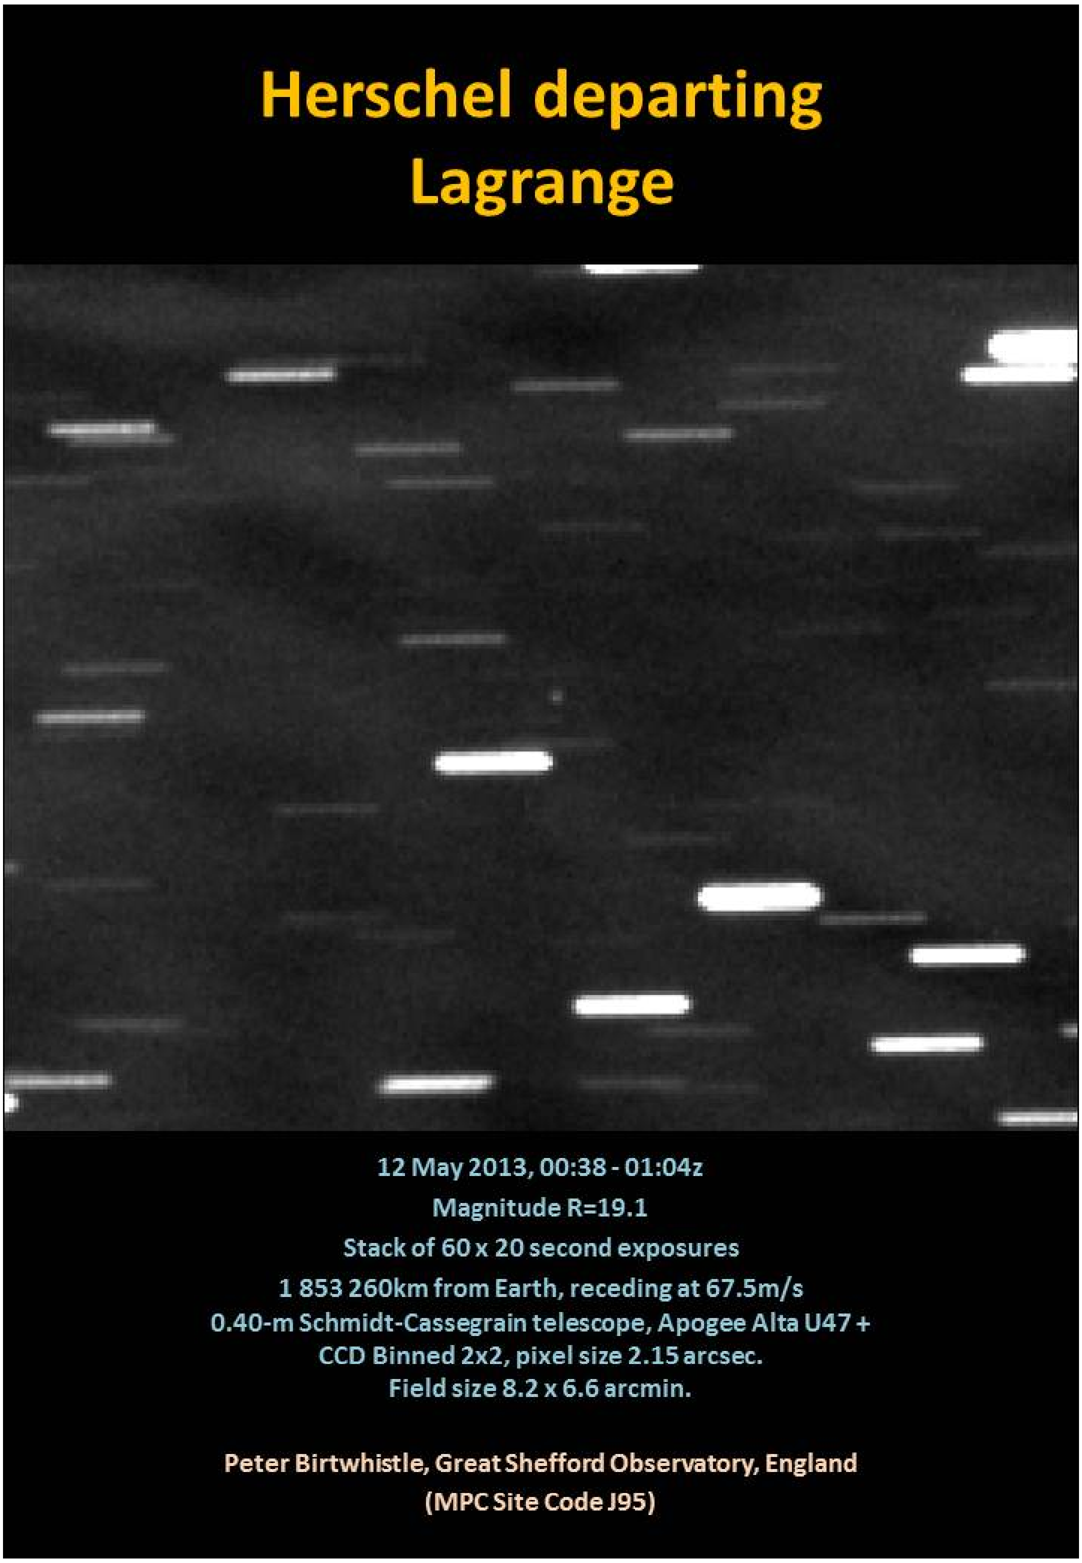

Figure 1.20. Herschel departing Lagrange, observed from Earth 17h before the main burn started. This image was taken on May 12th 2013, by British amateur astronomer Peter Birtwhistle, from Great Shefford Observatory (Minor Planet Center site code J93), approximately 90km west of London, using a 40cm reflector. The image is the sum of 60 individual exposures of 20s each, stacked on the motion of Herschel. The field of view is 6.6 arcminutes wide by 8.2 arcminutes high.

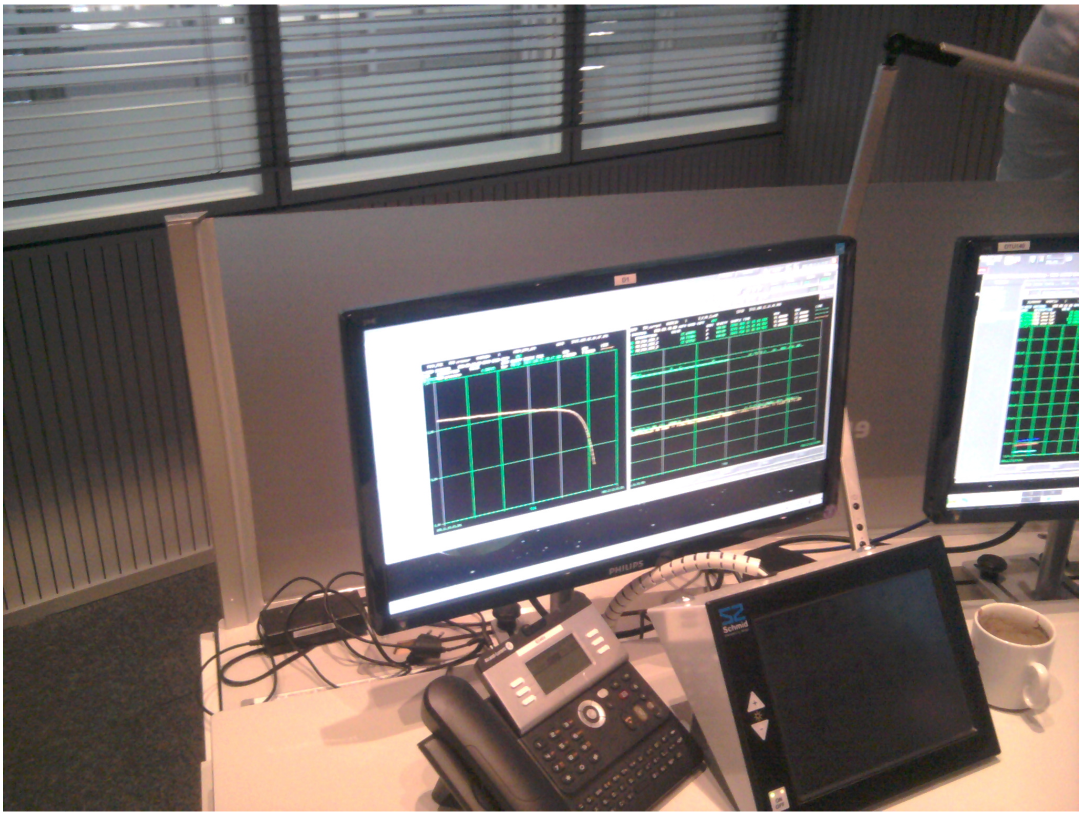

Figure 1.21. Screen in the Main Control Room at MOC showing (left hand plot) the rapid drop in pressure in the reaction chamber of the thrusters as the fuel was exhausted. This drop in pressure indicated that the fuel was close to exhaustion, allowing the final commands to cut the burn, stabilise the spacecraft and to passify to be carried out in a controlled manner.

Figure 1.22. Martin Kessler, Head of ESA's Science Operations Department, sends the final command to Herschel at 12:25 GMT (14:25 CEST), 17th June 2013, from the Main Control Room at ESOC, Darmstadt.

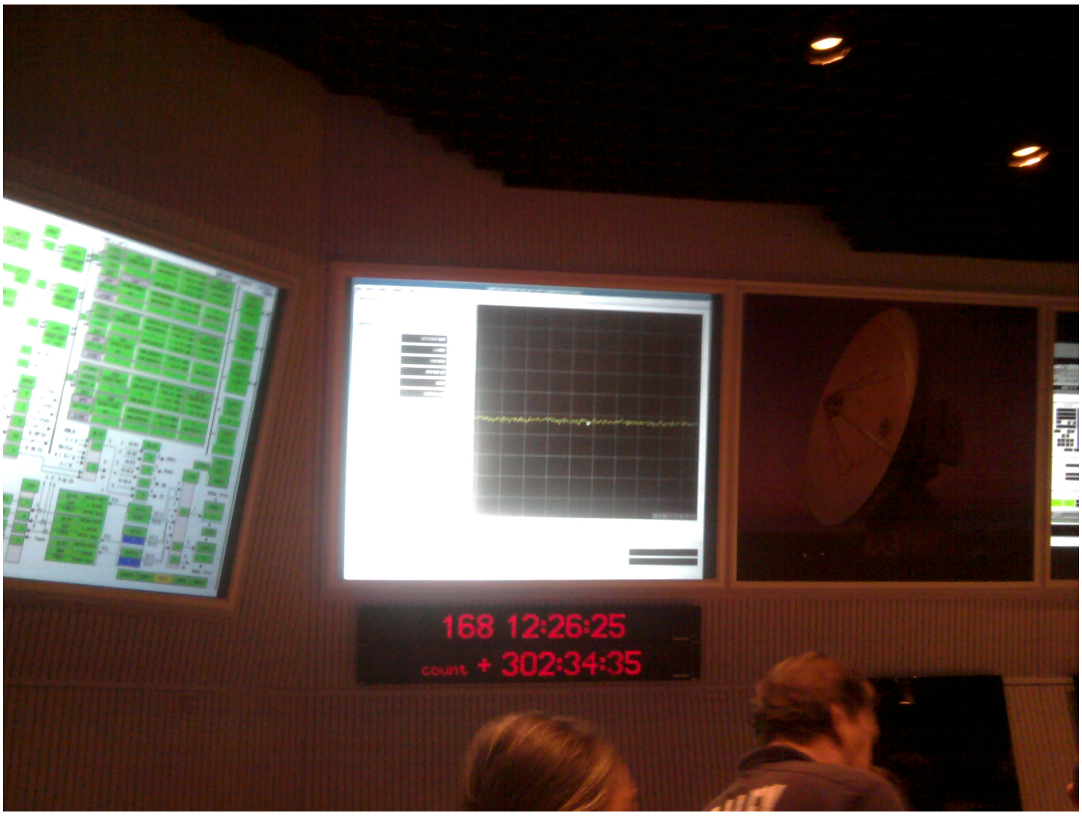

Figure 1.23. The screen in the Main Control Room at MOC showing the output from the frequency analyser. This would normally show a strong signal at the frequency of the Herschel transponder. The flat line shows that the signal from the Herschel transponder had dropped to zero after switch-off. The point at which the transponder signal dropped to zero marked the end of the Herschel mission.

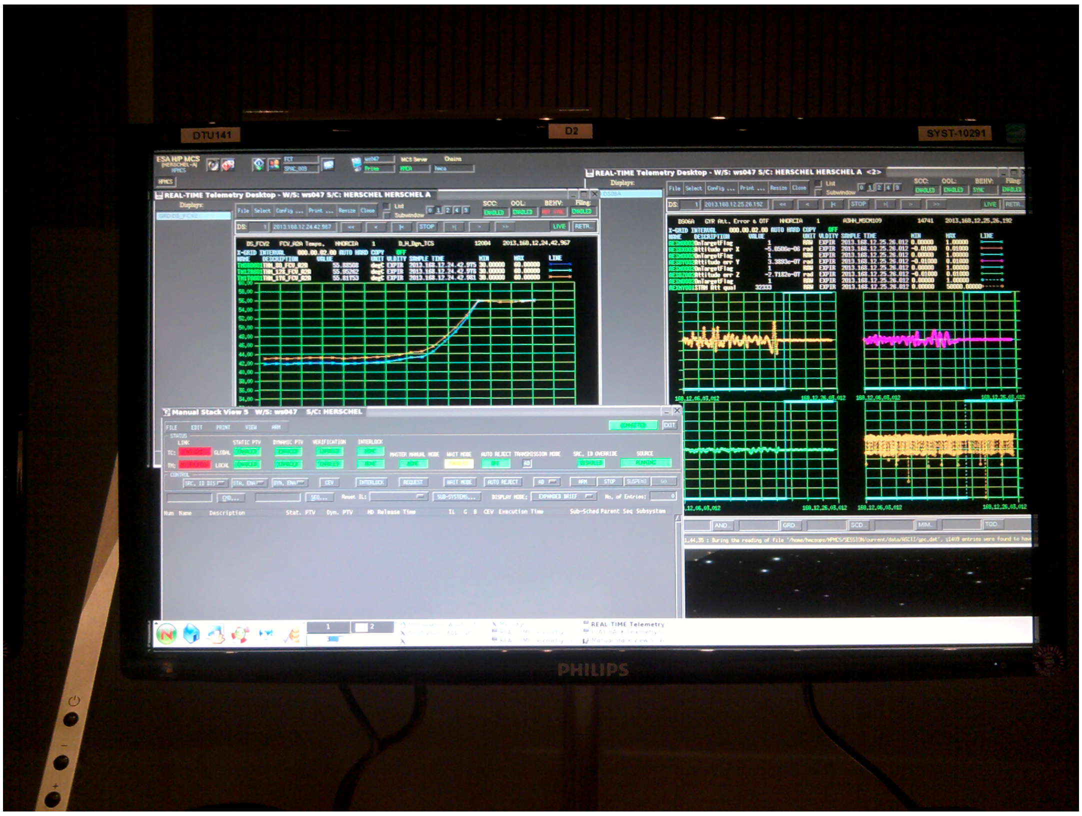

Figure 1.24. Telemetry on one of the consols in the Main Control Room at MOC showing the flatlining of the telemetry after transponder switch-off.

Astrometric observations of Herschel from groundbased observatories were made initially at four pairs of epochs after passivation, the last of them on 2013 July 1 (Figure 1.25). These revealed some small systematic residuals in the position of Herschel compared to the predictions made from the estimated impulse provided by the fuel drain burn. These were due to the difficulty in calculating the exact impulse given to Herschel by the drain manoeuvre with the pressure in the reaction chamber changing as the fuel was exhausted, given that no ranging measurements were available after the burn to give a definitive measurement. A correction derived from the astrometry was introduced into a refined orbit solution for Herschel. This found that the next Earth encounter would be at 8.236 million kilometres distance (21.4 Lunar Distances), significantly more than previously estimated, on October 13th 2027. The exact circumstances of this encounter have an exponential effect on the future evolution of the orbit, hence the interest in fixing the initial conditions as exactly as possible.

> **Note**
> As a measure of how sensitive the orbit of Herschel and the circumstances of the 2027 encounter are to initial conditions, the change in estimated encounter distance in 2027 from 19.5 to 21.4 Lunar Distances after the initial orbit update using astrometry to 2013 July 1, was due to a change of 41km in the adopted position of the spacecraft and 1.86m/s on its adopted orbital velocity on 2013 June 18.0.

The orbit is chaotic over centuries due to perturbations, in particular the cumulative effect of the regular Earth encounters. A small change in the circumstances of the 2027 encounter will have a disproportionate effect on the future orbit evolution. A second factor is that the effect of light pressure on the orbit is unquantifiable because it depends on the rotational state of the satellite once the reaction wheels have run down and no longer maintain it stable. Optical observations made from the ground in late June showed that the brightness of Herschel is varying on short time scales, showing that, as expected, it started tumbling after the reaction wheels had run down. However, orbit simulations made by MOC show that there can be no Earth impact for at least 300 years.

Due to the importance of determining the disposal orbit with the highest possible degree of accuracy, an attempt was made to recover Herschel before it became too faint and distant for groundbased observation. An initial attempt on 2013 September 9 was frustrated by cloud, but observations with the Fawkes South Telescope on 2013 September 10 were successful and a faint image of an object matching Herschel's motion, along the line of variation, obtained. The diference between the corrected and uncorrected ephemeris at this date was 5 arcminutes. Herschel was found to be 42 arcseconds from the prediction of the corrected orbit (corresponding to a 0.5 sigma deviation with respect to the revised orbit solution), allowing a further, significant increase in the accuracy of the orbit to be achieved by taking these observations into account in the orbit calculation. Confirmatory observations were taken with Pan-STARRS-w at Sutherland (South Africa), on September 11th; these are in excellent agreement with a rms residual for the five reported positions over the two nights of 0.32 arcseconds in Right Ascension and 0.38 arcseconds in Declination.

These data gave an approach to 0.057+/-0.002AU (8.5 million kilometres, or 22.2+/-0.8 Lunar Distances) on 2027 October 18, thus the tendency in orbit determinations has been to slightly delay the next Earth encounter and slightly increase the distance of its closest approach. However, the error on the prediction for the 2027 encounter was still extremely large: an error of 6 degrees along the line of variation.

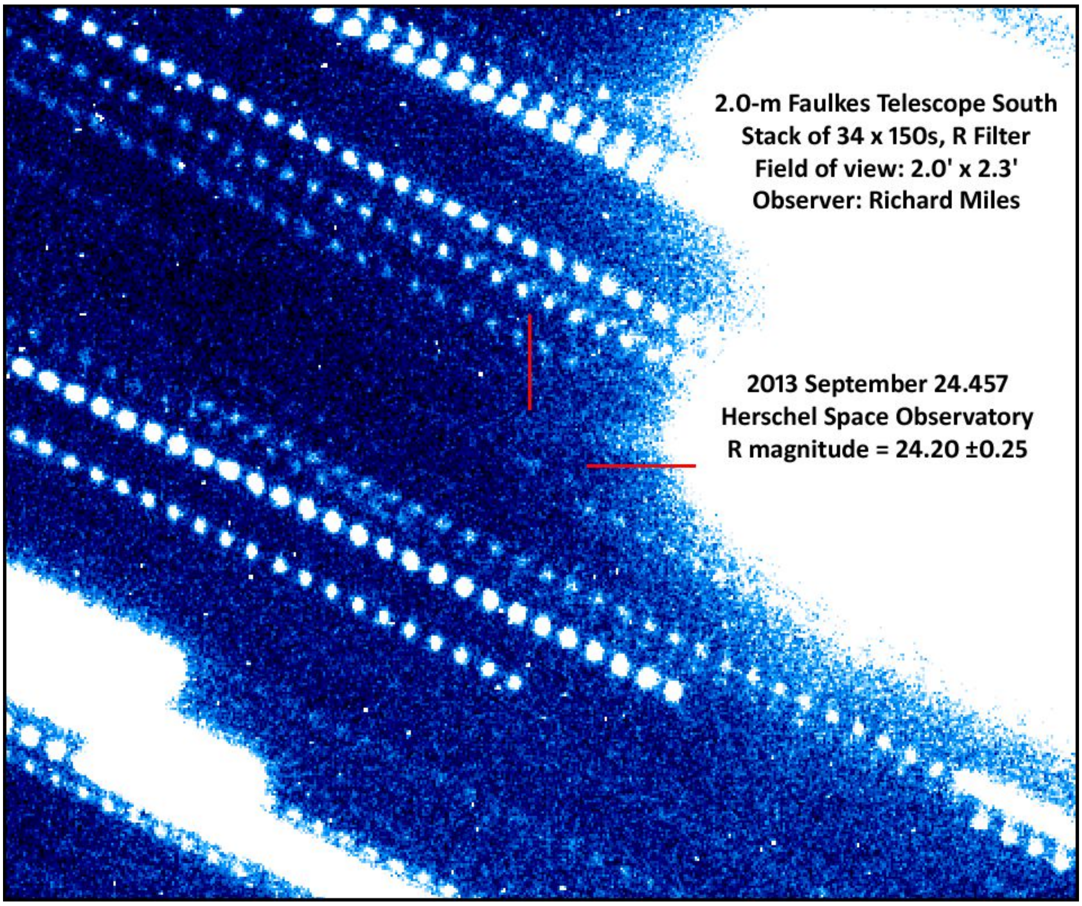

Figure 1.25. The last ground-based image of Herschel taken to date after passivation. This image was obtained on 2013 September 24 by British amateur astronomer Richard Miles, with the 2.0-m Fawkes South Telescope at Siding Spring. The total exposure is 85 minutes in very good seeing, made of 34 individual 150s exposures, tracked on the predicted motion of Herschel. Even in such a deep exposure, taken in good conditions, Herschel is right at the limit of the capabilities of this telescope.

Given the size of the uncertainties, a final attempt to acquire astrometry was made on 2013 September 24. At this date Herschel was at over 25 Lunar Distances and its predicted magnitude was below R=24. A sequence of 34 exposures of 150s each (total exposure 85 minutes), tracked on the predicted motion of Herschel were obtained in good seeing with the 2.0-m Fawkes South Telescope at Siding Spring Observatory. A very faint image was obtained at a magnitude of R=24.2 (Figure 1.25). This reduces the error in the predicted position for the 2027 return to 0.5 degrees, but the long-term orbit solution is limited by the (large) uncertainty in the light pressure term, as the rotational state of Herschel is completely unknown.

The final solution, using 15 observations taken on seven individual nights after the final drain burn, covering an arc from June 26 to September 24, gives an approach to 0.0557 AU on 2027 October 15, equivalent to 21.7 Lunar Distances.

The final orbit is very close to circular -- eccentricity, e=0.013 -- with semi-major axis a=1.042AU and a range from 1.0286-1.0558AU from the Sun. The period is 388.6 days.

Recovery of Herschel by ground-based observation at its next approach in 2027 would allow the long-term average of light pressure term to be calculated and would greatly improve the long-term orbit evolution determination. With the latest orbit solution it is now probable that Herschel can be deliberately recovered at its 2027 close approach by targeted observations and Herschel may even be bright enough to be picked up as a Near Earth Object with unusual motion by some automated surveys operating larger telescopes, such as PanSTARRS. Considerable efforts have been invested into ensuring that, should Herschel be detected in 2027, the orbit quality is good enough to ensure that it can be correctly be identified and that the orbit can be successfully linked to future serendipitous observations, so that it is not misclassified as a natural body in an Earth-threatening trajectory.

With no further ground-based observations reported by the end of September 2013, Herschel became too faint for further detection until its next close approach in 2027 and no further orbit refinement will be possible before then; at the end of 2013 the estimated magnitude of Herschel was around R=26, with the satellite now more than 0.25AU from Earth. The full light curve of Herschel from all ground-based observations from launch is shown in Figure 1.26; Herschel has been observed many times by automated asteroid search programmes searching around the opposition point, which have detected and reported it as an object with unusual motion in the sky. To date, 822 observations by ground based telescopes have been reported.

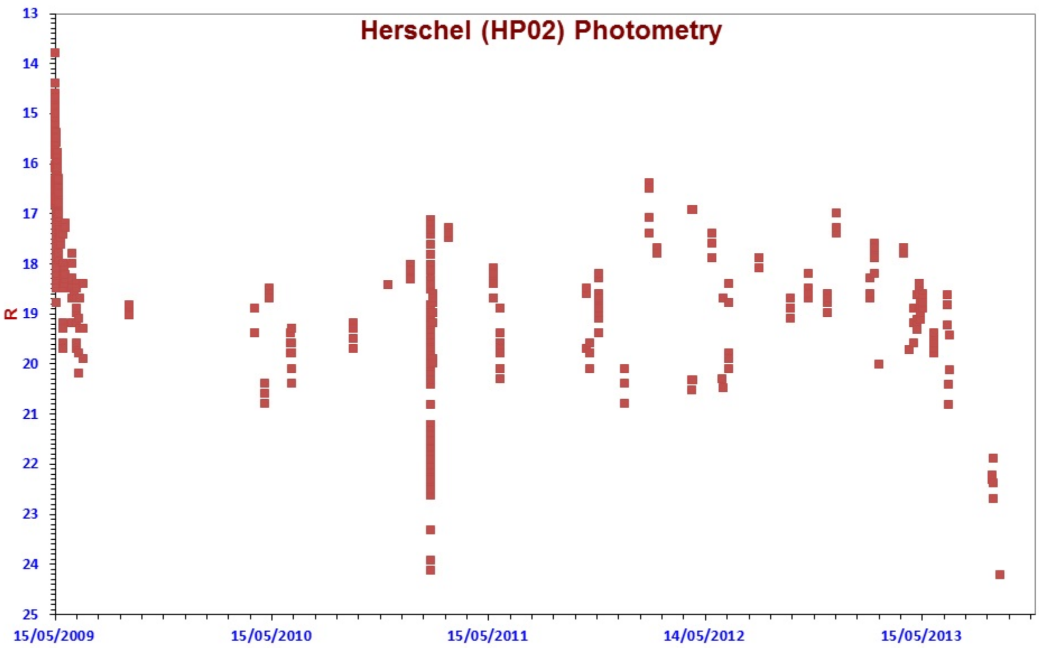

Figure 1.26. The light curve of Herschel from all observations reported to the Minor Planet Center (MPC) from launch until September 2013. During the initial stages after launch the different members of the Herschel-Planck constellation were reported as Planck, HP01 (Sylda), HP02 (Herschel), HP03 (Upper stage), HP04 (unidentified co-orbiter), etc. The brightness of Herschel was strongly dependent on its orientation and was, in a single night of observation, observed to vary by more than 5 magnitudes (greater than a factor of 100 in brightness) according to what reflecting surfaces were presented to Earth as the telescope slewed around the sky. The rapid fade as Herschel moved away from Lagrange in summer 2013 is obvious in this plot.

Other members of the Herschel-Planck constellation

Observers tracked various members of the Herschel-Planck constellation other than Herschel after launch. The ones with good orbit determinations are: Planck, HP01 (the Sylda) and HP03 (the upper stage of the Ariane 5); HP04 and HP05, both unidentified co-orbiters, have much lower quality orbits and were lost quickly. The main members of the Herschel-Planck constellation and their identification are listed in Table 1.2.

Table 1.2. Members of the Herschel-Planck constellation. Various alternative designations are found for individual members. Typically, in early mission phases, astrometry was reported using the Herschel-Planck (HP) constellation number, or the official launch designation (Year-sequential number). HP04 and HP05 were observed for only a few days and faded rapidly - it is quite likely that these were simply pieces of ice shaken off the cryogenic upper stage at separation. Various other, very faint, co-orbiting

fragments that were briefly reported from one or two observing stations were also, almost certainly, small fragments of ice that had separated from the upper stage.

Object	Official designation	Alternative designation
Herschel	2009-026A	HP02, NORAD 34937
Planck	2009-026B	NORAD 34938
Upper stage	2009-026C	HP03
Sylda	2009-026D	HP01
Unidentified	2009-026E	HP04
Unidentified	2009-026F	HP05

Planck made its short burn to leave L2 on August 14th 2013. On October 3rd the last science command will be sent, followed by the main burn on October 9th. The final command was sent on October 23rd 2013 after a drain burn on October 22nd 2013. The final orbit is similar to Herschel, although it had less remaining fuel and so has made a rather lower-energy escape. The orbit solution is available here.

HP01 was tracked from May 15th to August 15th 2009. 299 positions were reported. The Sylda made two orbits of the Earth, making an observed perigee pass in August 2009 and an unobserved one in October 2009, before escaping into heliocentric orbit. The orbit solution is available here. Bill Grey from ProjectPluto comments (private communication): "It looks as if it will return in 2025, but on the sunward side where we won't see it... that part is a little vaguer; the astrometry shows some pretty weird things, as if the object wasn't just moving under gravity and solar radiation pressure." The best solution is an approach to 0.052AU on 2025 April 20, but this is quite uncertain. Yarkovsky-like effects may be important in the orbit. Observations made in the nights after separation showed that the brightness of the Sylda was varying rapidly, showing that it was tumbling.

HP03 was tracked from May 16th to September 28th 2009. 353 positions were reported. The orbit solution is available here. Bill Grey from ProjectPluto comments (private communication): "The booster for Planck/Herschel will be back in March 2022, getting back up to mag 22.5 or so." The nominal solution is an approach to 0.092AU distance on 2022 March 30. As for the Sylda, Yarkovsky-like effects may be important in the orbit.

HP05 was tracked from May 16th to 24th 2009. 64 positions were reported from 4 observatories. The orbit solution is available here. Bill Grey from ProjectPluto comments (private communication): "It was clearly responding strongly to solar radiation pressure, i.e., it had to be pretty light. Based on its brightness, it was probably a meter or two across, but had a mass of under a kilogram. (There's some guesswork here, since we don't know the albedo of the object.)" The calculated mass to diameter ratio was exceptionally low for this object (7.2m^2/Kg) suggesting that it could be a sheet of ice shaken off the upper stage at separation.

HP04 was tracked from May 17th to 24th 2009. Only 32 positions were reported from 3 observatories. The orbit solution is available here. Bill Grey from ProjectPluto comments (private communication): "HPO4 had a similar brightness (to HP05) and, we assume, a similar size. There are some odd things in the astrometry that make me think it was manoeuvring a bit... maybe gas leaking out of the object." Again, its behaviour was consistent with a rapidly evaporating sheet of ice, which would have caused these asymmetric accelerations. This is supported by the fact that, as for HP05, its brightness dropped more rapidly than could be explained by its increasing distance from Earth.

A number of smaller and even fainter co-orbiting fragments were also detected briefly after launch. None had sufficient observations for a reliable orbit to be calculated and all faded away quickly, suggesting that they were nothing more than small pieces of evaporating ice shaken off the cryogenic upper stage.

1.2.4.2. Post-operations plans and objectives

Overall objectives

The goal of this phase is, within the constraints of time and available resources, to maximise the scientific return from the Herschel mission by facilitating continuing widespread effective and extensive exploitation of the Herschel data. This will continue after the conclusion of this phase (i.e. in the historical archive phase).

Documentation is being extensively revised during this phase to provide a legacy, special legacy documentation is being prepared and the Herschel Interactive Processing Environment (HIPE, see Section below) continues to be updated and refined in line with the state of knowledge of instrument behaviour and calibration and extra functionality is added progressively to allow users to obtain the best possible processing of their data.

The ultimate legacy of Herschel will be the historical archive, plus the sum of all the knowledge, both scientific and technical, derived from implementing and operating Herschel.

Documentation

One of the major objectives of post-operations is to consolidate Herschel knowledge. This means a complete update of all documentation to bring it up to date with the state of knowledge. In addition, there will be an Herschel explanatory supplement that will consolidate knowledge of both the mission in general and the individual instruments. The aim is to ensure that Herschel knowledge is conserved as a legacy for the future, long after post-operations have ended and the team has dispersed to other missions, or retired.

As an example of how much work and knowledge is embedded in the Herschel software, HSpot consists of approximately 160 000 lines of Java code, while the entire Herschel Common Science System, including HIPE and its documentation, represents approximately 4 million lines of code.

Data processing

Herschel's interactive processing environment is called HIPE. HIPE was used for all automatic overnight processing as data was received at HSC through the mission and has matured rapidly from its initial, pre-launch versions. HIPE 11 was released in late July 2013, while HIPE 12, built on Java 7, is expected to be released to users towards the end of 2013, or in early 2014. Up to and including HIPE 12, releases have come at six monthly intervals. After HIPE 12, releases will be yearly through to the end of the post-Operations phase, with the final release of HIPE expected to be HIPE 16. Users can find complete information about HIPE here.

During the post-Operations phase, refinement of the calibration of Herschel's instruments will continue as improvements are made in knowledge of the behaviour of the instruments, with more sophisticated processing and modelling. Although the "final" flux calibration at the end of operations was close to the definitive version, small updates and refinements will continue to be made in the future. These will be reflected in the bulk re-processings of the full Herschel Data Archive that will be made periodically when major new releases of HIPE become available, such that the products in the archive are the very best possible at any given time. These bulk re-processings will happen right through the entire duration of the post-Operations phase.

The Herschel data archive

Apart from ensuring the best possible legacy of Herschel documentation during this phase so that future users will be able to understand in detail how Herschel data was taken, what its peculiarities are and how it was processed, even when the relevant experts have moved on to other projects, or have retired completely from ESA, another fundamental objective is to ensure that the final data archive contains the best possible data products. Detailed information on the products in the Herschel data archive can be found in Section 8.1, or may be consulted on-line here.

To this end, all the Key Programmes - i.e. the large programmes which were awarded more than 100 hours of observing time by HOTAC - and the larger Open Time and Guaranteed Time programmes from the OT1 and OT2 Calls are invited to send their interactively processed data products to be incorporated in the Herschel Data Archive of User Provided Data. The User-provided data is of three basic kinds:

- Highly processed data (i.e. well beyond a quality obtainable with automatic pipeline processing).
- Published or publishable data (i.e. the final version of data that was used in publication(s), or that is to be published imminently).
- Data catalogues (i.e. catalogues of sources detected in the processed data, with its associated parameters).

Some of the user-provided data sets are very large. For example, the catalogue of extragalactic sources to be provided by the ATLAS consortium will contain approximately one million entries.

One of the features of user-provided data products is the fact that OBSIDs are linked to publications. In particular, it is possible to track what OBSIDs have been used in a particular publication and vice-versa. There is a link to an explanatory document published with the data set from each project, which may be a ReadMe file describing the dataset, or a data publication describing in detail how the data was obtained and processed.

All Herschel catalogues and observing logs will be available through VIZIER. As of early October 2013 the Herschel observing log is available for consultation in VIZIER.

The Herschel publication database

Another aim of post-Operations is to maintain a complete database of all publications made using Herschel data and link OBSIDs in the database, whether user-provided or not, to the publications in which they appear. Authors are strongly encouraged to list the OBSIDs that they use in publications to facilitate this tracking. As of December 2013 there were already 843 refereed papers published using Herschel data that are listed in the publication database, 809 of which (96%) had associated OBSIDs linked to the publication, while 34 more were yet to be examined at that point. This represented 34.7% of all the observing time awarded by HOTAC already being published in at least one refereed paper.

The publication database can be consulted here and will be updated constantly throughout post-operations.

1.2.5. Archive Phase

The historical archive phase is outside the funded Herschel mission. This phase commences after the end of the post-operations phase.

The historical archive will provide access to all Herschel observations and derived products. The products will all be derived in the archive consolidation phase, during the post-operations phase, in a consistent manner and to consistent standards, using the best knowledge of Herschel instrument calibration and data processing. In addition, the software, documentation -- manuals, etc.-- and tools will be available from the historical archive.

Chapter 2. The Observatory

This section summarises the main characteristics of the Herschel spacecraft, its orbit, pointing performance and observable sky regions.

2.1. Spacecraft overview

The Herschel spacecraft had a modular design, comprising the *Extended Payload Module (EPLM)* and the *Service Module (SVM)*. The EPLM consisted of the PLM "proper" with a superfluid helium cryostat - based on the proven ISO technology - housing the Herschel optical bench (HOB), with the instrument focal plane units (FPUs) and supporting the telescope, the sunshield/shade, and payload associated equipment. The SVM housed the "warm" payload electronics and provided the necessary "infrastructure" for the satellite such as power, attitude and orbit control, the onboard data handling and command execution, communications, and safety. The EPLM was based on the XMM-Newton Service Module. Figure 2.1 shows the main components of the Herschel S/C. Table 2.1 presents the Herschel Spacecraft key characteristics. The actual in-flight measured pointing performance is summarised in Table 2.4 in Section 2.4

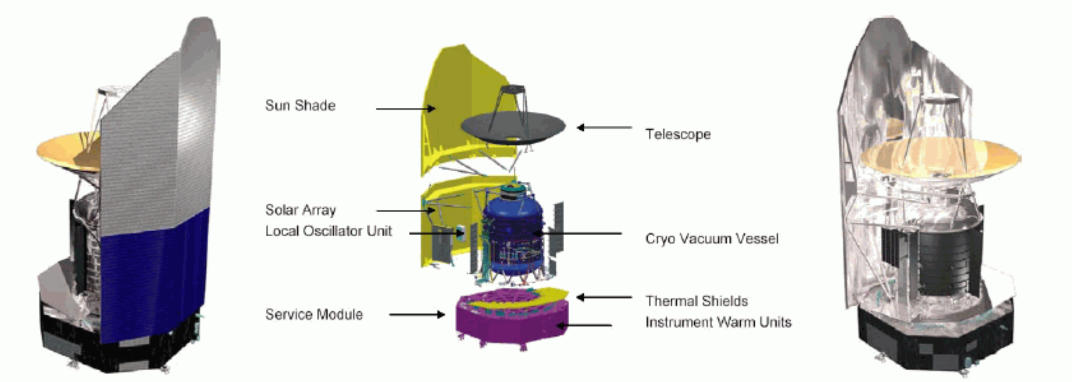

Figure 2.1. The Herschel spacecraft had a modular design. On the left, we see the "warm" side and on the right, we see the "cold" side of the spacecraft, the middle image details the major components.

Table 2.1. Herschel Spacecraft key characteristics

S/C Type:	Three-axis stabilised
Operation:	Autonomous (3 hours daily ground contact period)
Dimensions:	7.5 m high x 4.0 m diameter
Telescope diameter:	3.5 m
Total mass:	3170 kg
Solar array power:	1500 W
Average data rate to instruments:	130 kbps
Absolute pointing Error (APE):	1.90 arcsec (pointing) / 2.30 arcsec (scanning)
Relative Pointing Error (RPE, pointing stability):	0.19 arcsec (pointing)
Spatial Relative Pointing Error (SRPE):	< 1.5 arcsec
Cryogenic lifetime from launch:	min. 3.5 years (requirement), 3.96 years (achieved)

2.1.1. Herschel Extended Payload Module

The EPLM was mounted on top of the satellite bus, the service module (SVM) and consisted of the cryostat containing the instruments' focal plane units (FPU) and the Herschel telescope. The following sections describe the main components of the payload.

2.1.1.1. The Telescope

So that the favourable conditions offered by being in space could be exploited to the full, Herschel carried a precision, stable, low background telescope (Figure 2.2). The Herschel telescope was passively cooled by the large sunshade, allowing the size limitations imposed by active cooling using a helium jacket to be overcome. Thus its diameter was only limited by the size of the fairing on the Ariane 5-ECA rocket. The Herschel telescope had a total wavefront error (WFE) of less than 6 µm (corresponding to "diffraction-limited" operation at < 90 µm) during operations. It also had a low emissivity to minimise the background signal, and the whole optical chain was optimised for a high degree of straylight rejection. In space, the telescope cooled radiatively, protected by the fixed sunshade, to an operational temperature in the vicinity of 85 K, with a uniform and very slowly changing temperature distribution.

Figure 2.2. The Herschel telescope flight model seen in the clean room at ESTEC, prior to transport to Kourou.

The chosen optical design was a classical Cassegrain with a 3.5-m diameter primary and an "undersized" secondary. The telescope was constructed almost entirely of silicon carbide (SiC). The primary mirror (M1) was made out of 12 segments that were brazed together to form a monolithic mirror, which was machined and polished to the required thickness (~3-mm) and accuracy. The secondary mirror (M2), with 308-mm diameter, was manufactured in a single SiC piece. It was adjusted on the SiC barrel by tilt and focus adjustment shims. In order to avoid the Narcissus effect on the detectors, the central part of the secondary mirror was shaped in such a way that no parasitic reflected beam could enter the focal plane.

The hexapod structure (also made of SiC) supported M2 in a stable position with respect to M1. Finally, three quasi-isostatic bipods, made of titanium, supported the primary mirror and interfaced with the cryostat. The focus was approximately one metre below the vertex of M1, inside the cryostat.

The proper telescope alignment and optical performance were measured on the ground in cold conditions. The measured wavefront performance in cold was in line with the requirements. In-flight results confirmed the correctness of the focus position, although there was no possibility of in-flight adjustments such as focusing.

The M1 and M2 optical surfaces were coated with a reflective aluminium layer, covered by a thin protective "plasil" (silicon oxide) coating. The telescope was initially kept warm after launch into space to prevent it acting as a cold trap for outgassed volatiles while the rest of the spacecraft was cooling down.

Key telescope data are summarised in Table 2.2.

Table 2.2. The Herschel Telescope's predicted characteristics at a working temperature of 70 K.

Configuration:	Cassegrain telescope
M1 Free diameter:	3500-mm
Focal length:	28500-mm
f-number (M1/overall):	0.50/8.68
Field of View radius:	0.25°
M1 curvature radius / conic constant:	3499.02-mm / -1
Aperture stop / distance to M1 apex:	M2 mirror / 1587.555-mm
M2 diameter:	308.11-mm
M2 curvature radius / conic constant:	345.2-mm / -1.279
Image diameter:	246-mm
Image curvature radius / conic constant:	-165-mm / -1
On-axis best focus distance to M1 vertex:	1050-mm

2.1.1.2. The Cryostat

The Herschel cryostat housed the focal plane units of the three scientific instruments depicted in Figure 2.3. The cooling concept for the Herschel instruments was based on the proven principle used for the ISO mission. The temperature required in the instrument focal plane was provided down to 1.7K by a large superfluid helium dewar (helium at 1.6K), sized for a scientific mission of 3.5 years. This is achieved with a total amount of 2160 litres of helium cryogen. The cryostat provided 1.7K as its lowest service temperature to the instruments. Further cooling down to 0.3K, required for two instruments (the SPIRE and PACS bolometers), was achieved by dedicated ^{3}He sorption coolers that were part of the respective instrument focal plane unit. In orbit the liquid Helium was maintained inside the main tank by means of a phase separator (a sintered steel plug). The heat load on the tank evaporated the Helium over the mission lifetime at an estimated rate of about 200 grams per day. The enthalpy of the gas was used efficiently to cool parts of the instruments that did not require the low temperature of the tank (two temperature levels, at around 4K and around 10K). After leaving the instruments the evaporated gas was further used to cool the 3 thermal shields of the cryostat, before venting into space: the helium venting effectively counterbalanced almost exactly the light pressure on the sun shield.

During ground operations, the vacuum vessel was closed by the means of a cover, located at its top, which was opened once in orbit. To maintain a cold environment inside the cryostat during the last few days before launch in Kourou, an auxiliary liquid Helium tank was used. Once the fairing was closed before launch this auxiliary helium tank could no longer be used and boil-off began. The space side of the Cryostat Vacuum Vessel (CVV) was used as a radiator area to cool the CVV on orbit to a final equilibrium temperature of about 85K. This radiator area was coated with high emissive coating to achieve low temperatures in the L2 orbit. Multi-Layer-Insulation (MLI) covered the outer CVV-surfaces, in order to insulate it from the warm items (satellite bus and Sunshield). The outer layer of the MLI was optimised for the lowest temperature of the CVV. The outside of the cryostat was the mechanical and thermal mounting base for the Herschel telescope, the local oscillator unit of HIFI, the Bolometer Amplifier Unit of PACS and the large sunshield protecting the CVV from the sun.

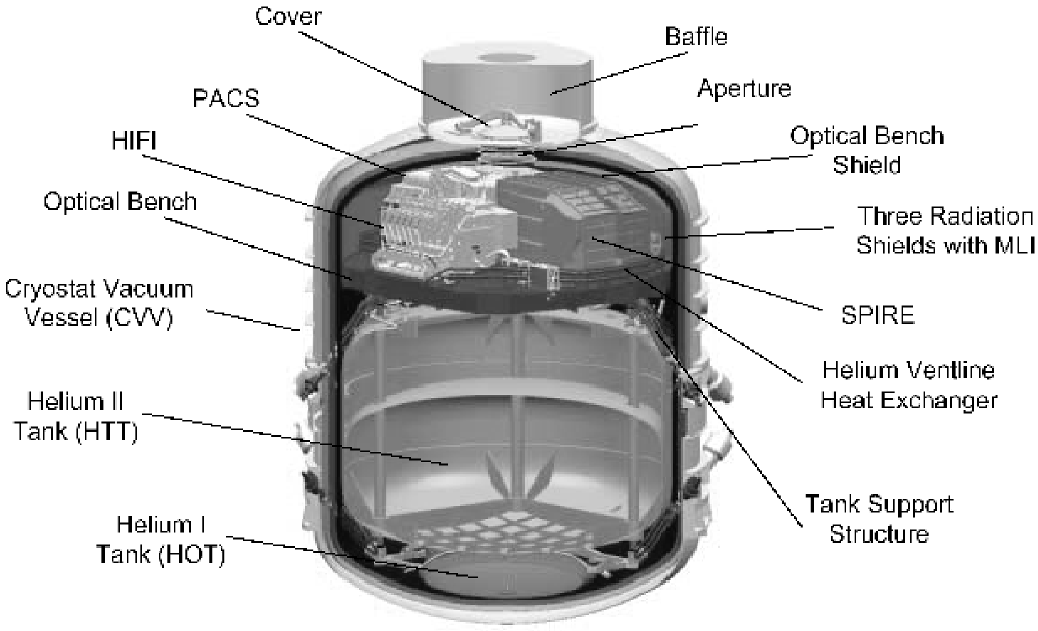

Figure 2.3. The Herschel cryostat.

2.1.1.3. Instruments

The science payload was accommodated both in the "cold" (CVV) and "warm" (SVM) parts of the satellite. The instrument FPUs were located in the "cold" part, inside the CVV mounted on the optical bench, which was sitting on top of the superfluid helium tank. They were provided with a range of interface temperatures from about 1.7 K by a direct connection to the liquid superfluid helium, and additionally to approximately 4 K and 10 K by connections to the helium gas produced by the boil-off of liquid helium, which was used efficiently to provide the thermal environment necessary for their proper functioning. The "warm" - mainly electronics - parts of the instruments were located in the SVM. The following instruments were provided within the Herschel spacecraft:

- The Photodetector Array Camera and Spectrometer (PACS)
- The Spectral and Photometric Imaging REceiver (SPIRE)
- The Heterodyne Instrument for the Far Infrared (HIFI)

The instruments are described in their respective users' manuals

2.1.2. The Service Module (SVM)

The service module (SVM) was the box-type enclosure at the bottom of the satellite, below the EPLM and carried all spacecraft electronics and those instrument units that operated in an ambient temperature environment. It is depicted in Figure 2.4.

SVM modularity was achieved by implementing units of similar function on each of the panels. Panels were either dedicated to one instrument, or to a single sub-system (Attitude Control, Power, Data handling-telecommunications). The propellant tanks were symmetrically placed inside the central cone. The SVM also ensured the mechanical link between the launcher adapter and the EPLM.

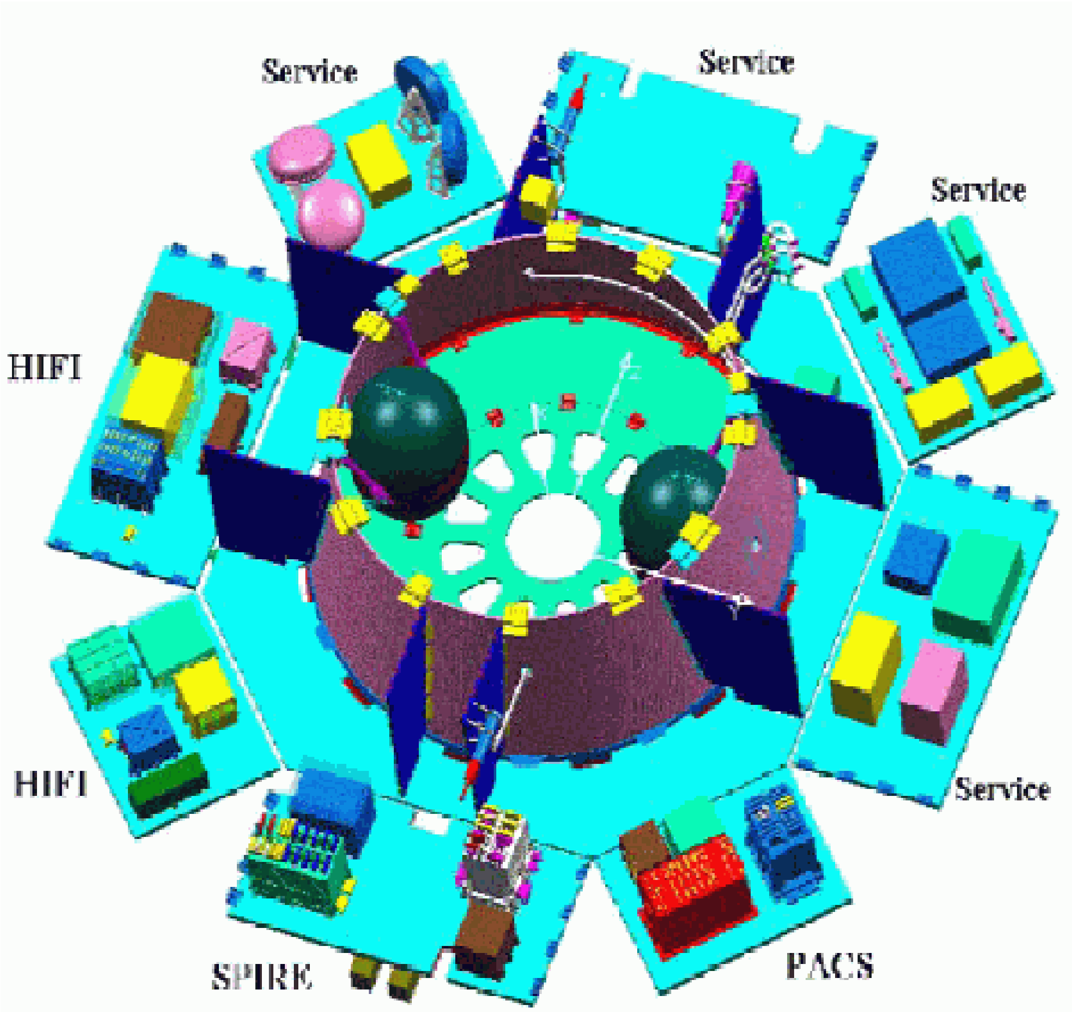

Figure 2.4. The Herschel service module.

2.1.2.1. The Sun shield and solar arrays.

The electrical power of the satellite was produced by the solar array. This was placed in front of the cryostat to protect the vacuum vessel from solar radiation. The rear of the sunshield was covered with multi layer insulation, as was the part of the cryostat facing this warm part of the system. The geometrical design had to consider the size of the cryostat and the telescope, the required solar aspect angles of the s/c in orbit and the limited diameter of the fairing of the launcher. For Herschel, a relatively simple system with a fixed solar array was selected. Only the lower part of the sunshield actually carried the solar cells. The upper part was free of solar cells to allow it to be at a lower temperature, which in turn helped the telescope to stay at the required temperature. The height of the sunshield was driven by the need to shade the entire telescope when the spacecraft was pointed closest to the sun (60° Sun aspect angle).

2.1.3. Spacecraft Axes definition.

The Herschel s/c coordinate axis system was defined in [RD1] as follows:

- The positive X-axis was perpendicular to the separation plane and nominally coincided with the longitudinal launcher axis. The positive X-axis was along the nominal optical axis of the Herschel telescope, towards the target source.
- The Z-axis formed a plane with the X-axis perpendicular to the separation plane such that nom-

inally the Sun lay in the XZ plane (zero roll angle), positive towards the Sun. In other words, the XZ plane was the plane of symmetry of the solar array, the Z-axis pointing outwards from the solar array.

- The Y-axis completed the right-handed orthogonal reference frame.

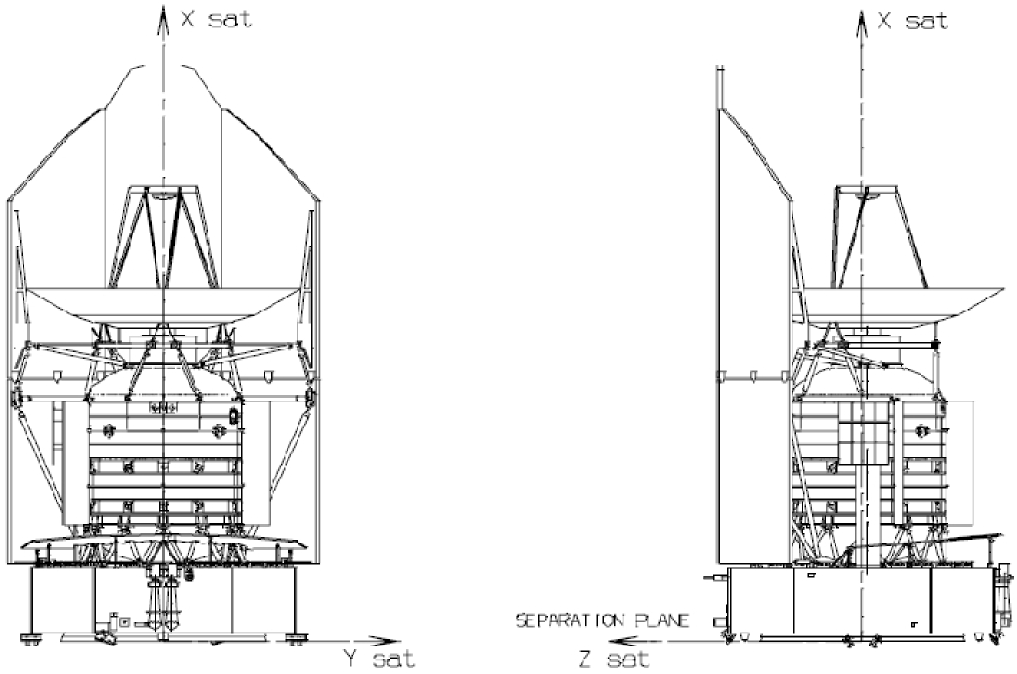

Figure 2.5. Herschel s/c axes (from [RD1])

2.2. Spacecraft orbit and operation

Herschel and Planck were launched aboard an Ariane 5-ECA launch vehicle from European spaceport at Kourou. The launch made use of the Sylda 5 adapter, with Planck being the lower passenger below the Sylda 5 and Herschel mounted as upper passenger. The two spacecraft separated within 30 minutes after launch and proceeded independently to different orbits about the second Lagrange point of the Sun-Earth system (L2) (see Figure 2.6).

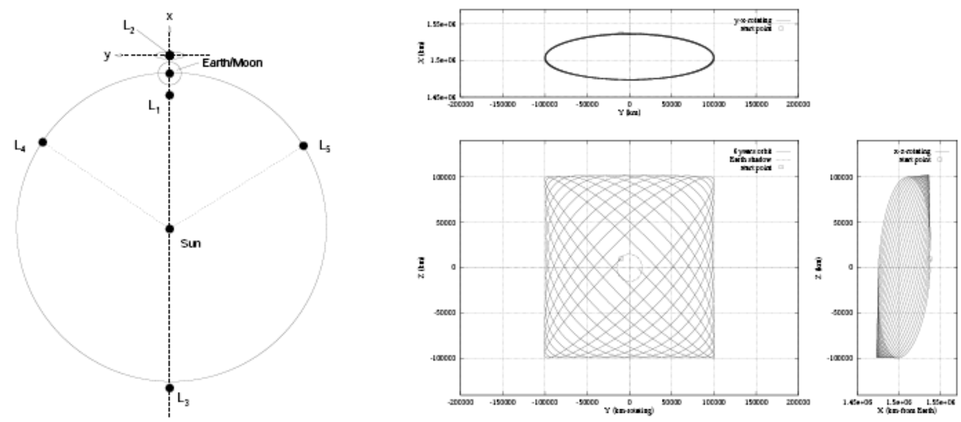

Figure 2.6. ***Left:*****The position of the Lagrange points for the Sun-Earth/Moon system. L2 lies 1.5 million kilometres from Earth.** ***Right:*** **An example of a Lissajous orbit around L2. The orbit x and y-axis are as shown in the plot on the left, the z-axis is normal to paper.**

Even though both satellites orbited around L2, their orbits were quite different: while Herschel went into a large, loose orbit that did not require an injection burn, Planck required a braking burn to go into its smaller and tighter orbit around L2. Herschel acquired its final orbital position at around 1.5 million km from the Earth after a transfer of about sixty days, with only a minor initial correction manoeuvre required early in cruise phase.

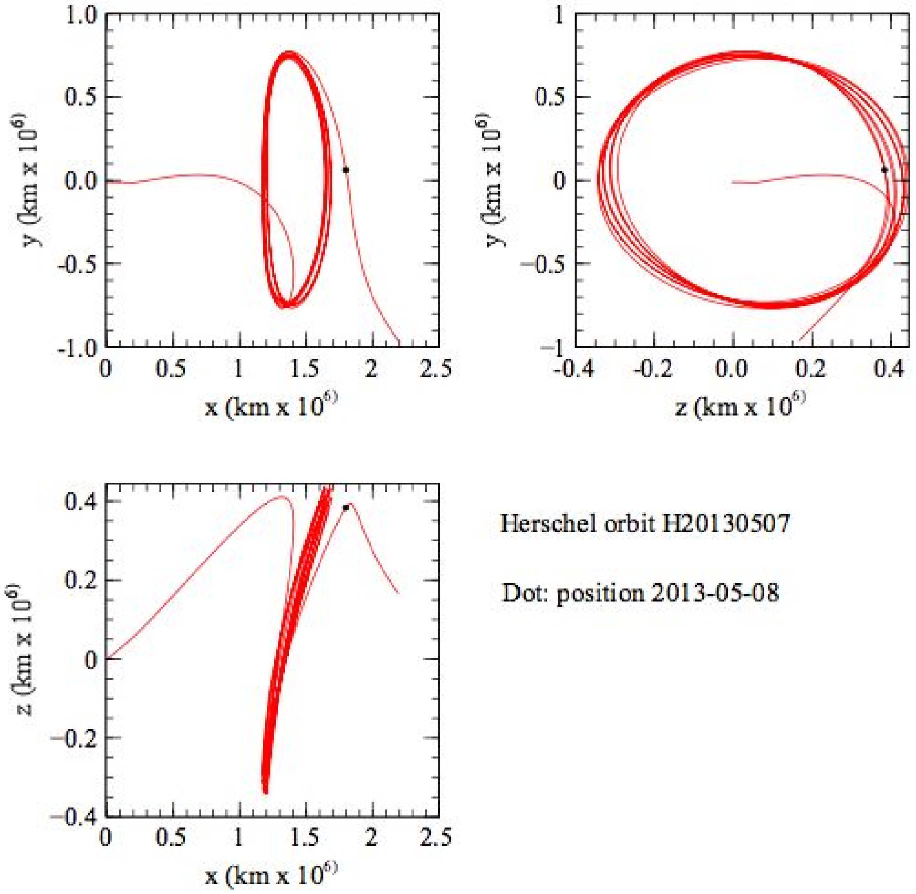

Figure 2.7. A representation of Herschel's orbit around L2 from top, side and front perspectives (i.e. seen along the z, y and x axes). The Earth is located at (0,0,0) in each plot. The red tracks are the projection on the three orthogonal planes of the 3D orbit from launch through to the end of June 2013, showing its position (blue dot) on May 8th 2013, nine days after End of Helium, already moving away from L2 as a result of the initial escape manoeuvre, which can be identified as a kink in the track, particularly in the perspective from above (top left), above the Earth's orbital plane. The horizontal and vertical axis scales are different, thus the orbit's shape is severely distorted in this view, although details of the individual loops around L2 are better visible than in an equal axis plot. A plot with the two axes on an identical scale to remove this distorion is shown in Figure 2.8. Plot generated by Jon Brumfitt based in the orbit file generated by the Flight Dynamics Team at MOC.

The Herschel spacecraft was eventually placed in a large "halo" orbit around L2 (halo orbits are special cases of Lissajous orbits around Lagrange points where the in-plane and out-of-plane frequencies are the same), with an amplitude of about 700 000km and a period of approximately 178 days. The distance from the Earth ranged from 1.2 to 1.8 million km.

The orbit chosen for Herschel presented a number of advantages summarised below:

- It simplified long observations, since the Sun and the Earth remained close to each other as seen by the S/C (Sun-S/C-Earth angle was always < 40°)
- It gave a very stable thermal and radiation environment
- There was no trace of atmosphere
- The environment at L2 is clear of space debris
- A large halo orbit can be achieved without an injection Δv

Major drawbacks were the long distance for communications, requiring a 35-m class antenna and the fact that orbits around the L2 are unstable; without orbit corrections the spacecraft would have deviated exponentially from the nominal one. Small correction manoeuvres, applied at approximately monthly intervals, maintained the orbit close to the nominal one. Figure 2.7 shows the actual Herschel halo orbit around L2.

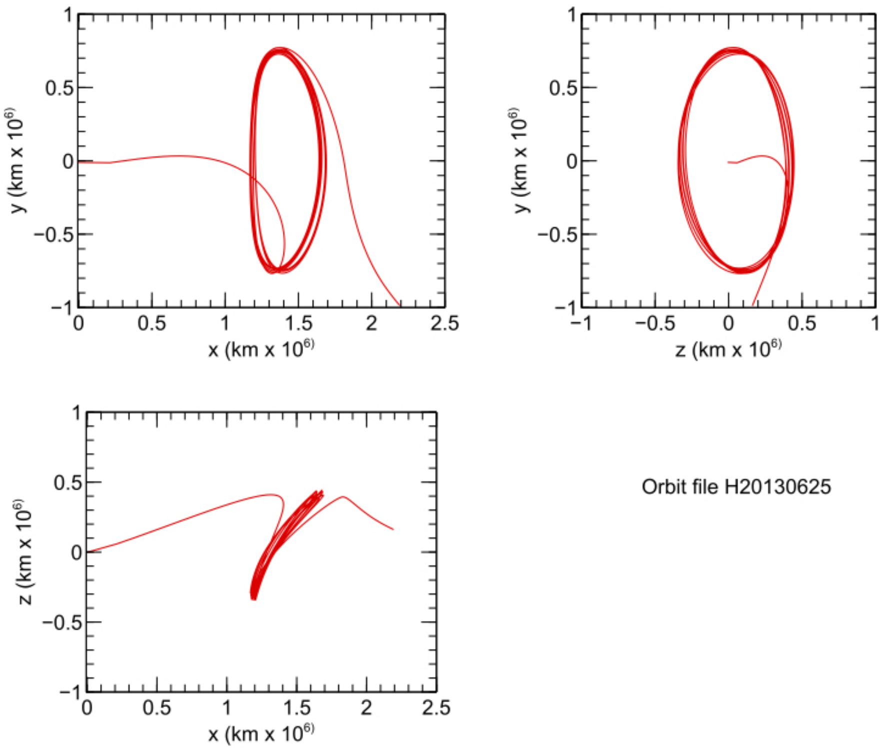

Figure 2.8. Herschel's orbit from launch to the end of June 2013, as for Figure 2.7 plotted with equal axis scales to give a more faithful representation of the orbit shape, at the cost of some loss of fine detail. The rapidity of the satellite's drift away from L2 after the disposal burn is manifest in this image. Plot generated by Jon Brumfitt based in the orbit file generated by the Flight Dynamics Team at MOC.

Herschel operations were performed by the European Space Operations Centre (ESOC) located in

Darmstadt (Germany). The main ground station was New Norcia (Australia), which is equipped with a a 35-metre antenna using X band up and down links. New Norcia was backed up by the Cebreros ground station (Spain). In the phase immediately after launch the Kourou (French Guiana) was also used. During routine operations, the ground station communication link was restricted to a duration of approximately 3 hours. During this time, the spacecraft antenna was pointed to the Earth. The data stored in the on-board solid state mass memory were downlinked, and the mission time line with the new schedule was uplinked. Real time operations and spacecraft maintenance were also carried out during this period. The rest of the time the satellite operated autonomously. The system was designed to support 48 hours of autonomous operation, with required a solid state mass memory capability of 25 Gbt. The amount of Herschel data downloaded per day was in excess of 8 Gbt.

2.3. Sky visibility

The areas of the sky accessible to the Herschel telescope were determined by a number of constraints applicable to Sun, Earth, Moon and other bright solar system objects. In particular, the following constraints were applicable through the mission:

- Sun-S/C-LoS angle in the S/C XZ plane (Solar Aspect Angle or SAA) of 60°.8 to 110° for normal operations. The allowed range was reduced with respect to the original one (60° to 120°) as, post launch, a small change was made to the extreme range of permitted of solar aspect angles to limit it to a maximum range from 60°.8 to 119°.2.
  .

  Similarly, since in the extreme SAA range ('warm' attitude range, SAA in the 110° to 119°.2 interval), a noticeable pointing performance degradation (larger APE and pointing offset drift) due to thermo-elastic effects was observed. Moreover, this degradation persisted even when the S/C was brought back to 'cold' attitude until the structure settled back in the original position. As the telescope pointing performance improved it was noticed that the degradation was present at even smaller SAAs from 105-110°, leading to a further restriction in the freely scheduled "cold" area of the sky.

  Nevertheless, if deemed unavoidable, short (normally less than 1 hour) observations in the 'warm' SAA range (105° to 119°.2) could be scheduled, normally at the end of operational days. The observer needed to be aware that in such cases, a degradation of the pointing accuracy was very likely, particularly beyond 110°. Longer observations would only be scheduled if scientifically strongly justified and taking care to minimise the impact on observations from other programmes. On some occasions it was necessary to hold Herschel at hot attitudes for as long as two days for instrument recovery activities: in such cases only observations robust against pointing drift would be scheduled until cold conditions had been restored.

  Similarly,

- Maximum roll angle of ±1°

In addition, the following extreme Earth and Moon angles occurred across the mission (to be taken into account for straylight considerations):

- Sun-S/C-Earth angle of 37°
- Sun-S/C-Moon angle of 47°

In order to avoid straylight pollution and also for safety reasons (to prevent large fluxes of light from reaching detectors), the nominal half-cone exclusion angles listed in Table 2.3 applied to observations towards major planets.

Table 2.3. Nominal exclusion angles (half-cones) for observation towards major planets

Instrument	**Mode**	**Mars**	**Jupiter**	**Saturn**	**Instrument Critical**
SPIRE	Slew	15 arcmin	15 arcmin	15 arcmin	Yes[a]
	Pointing	1.5 deg	1.5 deg	1.5 deg	Yes[a]
HIFI	Slew	36 arcmin	36 arcmin	36 arcmin	No[c]
	Pointing	36 arcmin	36 arcmin	36 arcmin	No
PACS	Slew[d]	4 arcmin	4 arcmin	4 arcmin	No
	Pointing[e]	1.5 deg	1.5 deg	1.5 deg	No

a. SPIRE determined that, while Jupiter and possibly Saturn would not damage the instrument, they would have rendered it inoperable for a significant period (possibly even an entire OD)

b. For SPIRE PACS parallel mode both the SPIRE and PACS restrictions applied.

c. HIFI wished to avoid straylight pollution when observing fainter objects with a SSO close to the instrument LoS. The instrument would not be harmed by the presence of a major SSO in the FoV and did, in fact, even use Mars as its primary calibrator.

d. During slews, the detectors were ON (photometry, spectroscopy or parallel mode).

e. During non-SSO PACS observations. PACS also observed these SSOs directly.

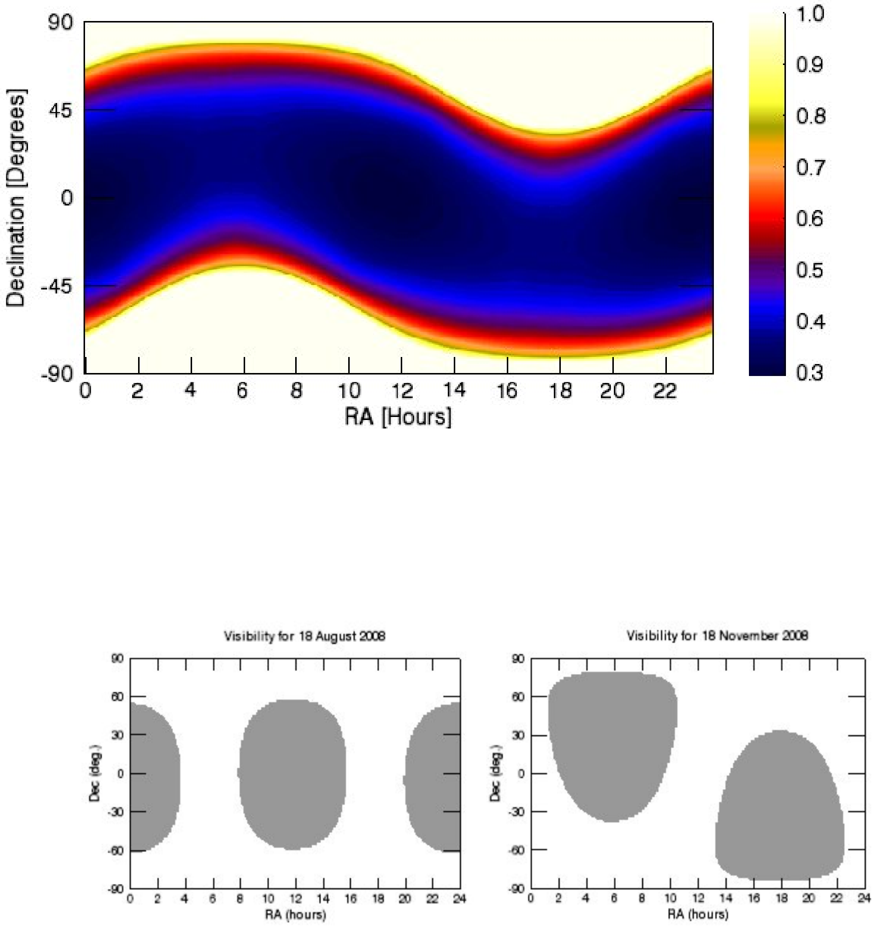

Figure 2.9. Top: The sky visibility across the sky as a fraction of the total hours through the Herschel mission, represented as a colour scale (shown at right) where black represents 30% visibility and white

represents permanent sky visibility. Bottom: sky visibility for two sample dates. Shadowed areas represent inaccessible sky areas.

The time windows when a fixed or moving target or list of targets were visible can be calculated with HSpot. The tool provides an easy way to check in which time intervals a source was visible during the mission. The visibility calculation does not take into account the avoidance cones around Jupiter, Saturn and Mars described above.

The sky visibility for each date has been determined by the launch date (14th May, 2009) and the orbit of the satellite. Considering the nominal duration of the operations, all areas in the sky were visible at least 30% of the time. The sky visibility region moved slowly on a daily basis. The two snapshots at the bottom of Figure 2.9 illustrate the typical sky visibility differences after a 3 month interval; although this was calculated for an different launch date to the actual one, the graphic remains a valid representation of the effect.

2.4. Herschel pointing performance

This section deals with the pointing performance of the Herschel spacecraft. The spacecraft Attitude Control and Measurement System (ACMS) consisted of several components, as depicted in Figure 2.10. The main constituents of the ACMS were the attitude control computer (ACC), gyroscopes (GYR), star trackers (STR), reaction control system (RCS), reaction wheel assembly (RWA), Sun acquisition sensors (SAS), coarse rate sensors (CRS) and attitude anomaly detectors (AAS).

2.4.1. Attitude control and the Startracker system

In normal operation, the spacecraft attitude was commanded by means of the reaction wheel system. It comprised of four 8.6 kg wheels in a skewed configuration, each with a momentum storage capacity of 30 Nms and a maximum delivered reaction torque of 0.215 Nm in either the positive or the negative direction.In the baseline configuration, all four whels are powered and used for actuation, providing optimum slew performance and momentum storage. Nevertheless, the ACMS was also capable of operating with only three reaction wheels powered. In the nominal configuration, the maximum slew speed was 0.00204 rad/sec, i.e. ~ 7 arcmin/sec.

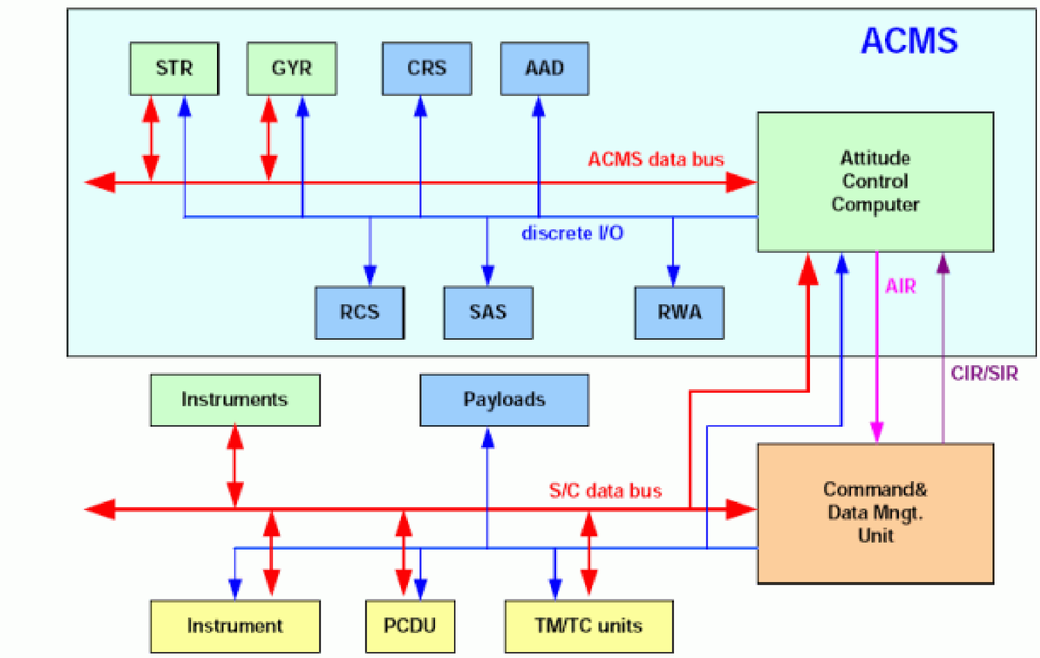

Figure 2.10. A schematic diagram of the Herschel/Planck avionics.

In normal science operation, the spacecraft attitude was controlled by means of two components: the star trackers (STR) and gyroscopes (GYR). The STR comprised of two cold-redundant units, nominally aligned with the -X axis. The STR hardware included:

- An objective lens.
- A baffle to protect from undesired straylight from the Sun and other bright sources.
- The focal plane assembly, containing a CCD detector and a thermo-electric cooler for CCD cooling.
- The sensor electronics.

From a functional point of view, the STR can be seen as a video camera plus an image processing unit that, starting from an image of the sky, extracted the attitude information measured with respect to the J2000 inertial reference system and delivered it to the ACC. A CPU (ERC32 microprocessor) controled the CCD sensor and also carried the image processing task.

The main characteristics of the Herschel's STR were:

- The ability to determine the inertial position from "lost in space".
- FoV: 16.4×16.4 deg^2.
- An onboard catalogue, based on Hipparcos, of some 3000 bright stars.
- A minimum of 3 stars, 9 was the maximum due to H/W limitations.

The STR bias was the largest contributor to the absolute pointing error and is pixel-dependent (some 0.8" $\times \sqrt{2}$)

The STR was provided with an enhanced performance mode the so-called "interlaced mode", only applicable if there were ≥ 15 stars in the FoV. The STR sampled at twice the nominal frequency (4 Hz), measuring 9 stars at a time.

In order to get the maximum accuracy it was necessary for the ACC to provide an accurate value of the S/C angular rate as input to the STR (the maximum performance was achieved with rate errors below 0.2 arcsec/sec).

Gyroscopes (GYR) are devices that use a rapidly spinning mass to sense and respond to changes in the inertial orientation of it spin axis. *Rate/rate-integrating gyros* provided high-precision measures of the the spacecraft angular rate. The Herschel's ACMS was provided with four gyroscopes mounted in a tetrahedral configuration. The gyroscopes were hot-redundant, and each of the four could replace any of the others. The fourth gyroscope was not used for control, but served to detect an inconsistency in the output of the other three.

The STRs provided an absolute reference, but with limited accuracy. In contrast, the GYRs were very accurate, but only on short temporal (bias drift, 0.0016 deg/hour) and spatial (variation in the scale factor should be taken into account for distances larger than 4 deg) scales. Therefore, the GYR attitude must be recalibrated using the STR information. As a result, in normal operation the spacecraft attitude was computed by combining the STR and GYR measurements in the ACC using a linear Kalman filter. The so-called "filtered attitude" was sampled and downloaded with a frequency of 4Hz.

2.4.2. The Startracker catalogue

The Startracker used a catalogue of 3599 stars measured by Hipparcos, corrected for proper motion to the mid-date of the mission. Of these, 203 stars were blends of two stars, resolved by Hipparcos, but not by the Herschel Startracker. This required a barycentre to be calculated of the two Hipparcos components to give the Herschel position, taking into account the Herschel instrumental magnitudes. Finally, in analysis of tracking performance, 72 stars were found to give systematic residuals in pointing and were suppressed from the fine attitude guidance catalogue from OD-1032 [RD12].

2.4.3. Herschel pointing modes

Herschel pointing modes were based either on either stare pointings (fine pointing mode), or moving pointings at constant rate (line scan mode). These come in multiple flavours:

- Raster maps are 'grids' of stare pointings at regular spacings
- In the position switching and nodding modes, the boresight switched repeatedly between two positions in the sky.
- Scan maps are sequences of line scans at regular spacing. Allowed angular speed ranged from 0.1 arcsec/sec to 1 arcmin/sec.
- In addition, the Herschel spacecraft could track moving Solar System targets at rates up to 10 arcsec/min, with faster rates up to 30 arcsec/min possible, at a risk of a degraded tracking performance although, in practice, no such degrading was seen.

2.4.4. Pointing accuracy definitions

The formal definitions of the spacecraft pointing accuracy parameters are provided in this section. Five different measures of pointing accuracy are commonly used and are each defined below.

The term 'pointing', when applied to a single axis (e.g. the telescope boresight), refers to the unambiguous definition of the orientation of this axis in a given reference frame. When characterising the pointing performance of the telescope, it is possible to provide a figure of the absolute attitude accuracy provided by the ACMS (absolute pointing error), or how accurate the 'a posteriori' knowledge of the absolute attitude (the absolute measurement error) can be, or how stable the pointing is (the relative pointing error). Furthermore, the pointing performance can be also characterised in terms of the relative accuracy of a set of attitude measurements (the spatial relative pointing error). The latter measurement is important to characterise the accuracy of the relative astrometry in a map comprising several pointings (e.g. from a raster pointing).

Herschel pointing accuracy definitions, presented below, are based on the prescriptions given in the *ESA Pointing Error Handbook* (ESA-NCR-502):

- **Absolute Pointing Error (APE)**: the angular separation between the desired direction and the actual instantaneous direction.
- **Absolute Measurement Error (AME)**: the angular separation between the actual and the estimated pointing direction (*a posteriori* knowledge).
- **Pointing Drift Error (PDE)**: the angular separation between the average pointing direction over some interval and a similar average at a later time.
- **Relative Pointing error (RPE) or pointing stability**: the angular separation between the instantaneous pointing direction and the short-time average pointing direction during a given time period (in this case 60 sec).
- **Spatial Relative Pointing Error (SRPE)**: angular separation between the average orientation of the satellite fixed axis and a pointing reference axis, which is defined to an initial reference direction.

2.4.5. The pointing performance of Herschel and its refinement

2.4.5.1. Pointing performance summary and pointing error sources

The main pointing error contributors within the Herschel spacecraft were:

- To AME and APE:
  - Position-dependent bias within STR. *It is also the main contributor to SRPE.*
  - Residuals from calibration
  - Thermo-elastic stability of the structural path between STR and FPU
  - Instrument LoS calibration accuracy w.r.t. ACA frame (best for PACS)
- To PDE: Thermo-elastic stability
- To RPE: The main contributor is the noise in the control loop, comprising STR+Gyro noise attenuated by a linear Kalman filtering.

Table 2.4 summarises the pointing performance of the Herschel spacecraft.

Table 2.4. Herschel pointing requirements (from SRS v3.2) compared with predictions and measured performance. Goal conditions assume that 18 stars were available for guidance within the STR field, allowing interlacing to be performed.

	Baseline (arcsec)		**Goals (arcsec)**	
Name	Requirement	Performance	Requirement	Performance
		Predic./Measur.		Predic./Measur.
APE point	3.7	2.45/0.90	1.5	1.45/0.8
APE scan	3.7	2.54/0.8-0.9	1.5	1.63/n.a.
SRPE	1.00	2.44/1.0*	1.00	1.52/1.0

*There could be a marginal non-compliance in the SRPE as this figure was measured in very small raster maps and was quite uncertain.

2.4.5.2. Pointing performance refinement

There was a considerable improvement in the pointing performance during the mission as the Startracker system was better understood and characterised. This was attained in a series of stages, with the impact of each one being carefully assessed before the next step was taken. This was a result of the considerable efforts of many people at MOC, HSC and in industry[RD12].

In the very earliest mission phases after launch, before the cryocover was opened, it was found in initial tests that the heaters for the gyros were causing magnetic fields that created an oscillation in the pointing. This was cured by turning off the heaters, which could be done without the gyros cooling beyond acceptable temperatures. Similarly, the Startracker CCD performance was significantly improved by lowering its operating temperature, on OD-320, from the initial +20°C to -10°C to reduce the number of bad pixels and, in combination with regular characterisation and update of the on-board bad pixel table, this eliminated the so-called "speed bumps" where false detections of stars in the Startracker field caused jumps in scans.

An important factor in the first year of the mission was to characterise the SIAM (Spacecraft Instrument Alignment Matrix) file. This matrix told the spacecraft where the instrument apertures were in relation to the centre of the field of view. The first SIAM was used on OD-38, using Sneak Preview data (see (Section 1.1.2), applying PACS 70 micron data to provide an initial correction to the pre-launch SIAM. Further updates were made as individual instruments obtained and analysed observations that allowed them to characterise the different instruments. These initial determinations were applied by PACS on OD-53, HIFI on OD-55 and SPIRE on OD-68. Further refinements were applied on OD-78, OD-82, OD-98 and OD-122, as additional observations allowed the initial estimates to be improved. The OD-122 update was the one used for most of the first year of operations. As the amount of observational data increased, with massive scheduling of science data, first in SDP and then in routine phase, it was recognised that a further, small adjustment was required to the SIAM. This was applied on OD-341 and was the last update required during the mission.

A few observations in certain sky directions were found to show systematic pointing offsets that were found, on investigation, to be due to a strong asymmetry in the distribution of the guide stars, with all, or almost all, being on the same side of the CCD field. These offsets could be understood as due to small distortions in the focal plane of the Startracker. These focal plane distortions were corrected in three stages.

The first stage was to introduce a focal length correction. This was performed on OD-762 and led to an improvement in the APE from 2.5 arcseconds to 1.45 arcseconds.

Further analysis showed that the introduction of a one-dimensional distortion correction would give rise to a further substantial improvement in the pointing. This was applied on OD-866 and led to a further improvement of the APE to 0.95 arcseconds; considerably less than the pre-launch goal for pointing.

With further study, it was realised that second and higher order polynomial distortions were still present. It was thus decided to attempt a further, more ambitious two-stage refinement of the pointing: the first stage was to apply a n-dimensional distortion correction; the second to apply a clean-up of the guide star catalogue, removing those stars that were found to give systematically larger residuals (this was the suppression for use in fine pointing of 2% of the stars in the catalogue, which were found to be consistent outliers). To complete the first part, it was necessary to determine 16 coefficients. This was done from five independent data sets, each constructed from twenty ODs of tracking data, which gave consistent solutions for the coefficients.

The first part was implemented on OD-1005 and, on being verified as being correctly implemented, was made definitive on OD-1011. The second part, the catalogue clean-up, was applied on OD-1031, was the final fine tuning. The effects of these improvements were checked with telescope pointing tests on OD-1005, OD-1028 and OD-1034 that verified further substantial improvements in pointing performance to the final values seen in see Table 2.4; the final, measured APE after both corrections had been applied was 0.81 arcseconds. These improvements also showed that the SIAM that had been applied on OD-341 required no further adjustments.

Chapter 3. Overview of scientific capabilities

Herschel was a versatile observatory with a wide range of capabilities that covered point-source photometry, imaging, large area mapping and spectroscopy at both intermediate and high resolution. Despite the relatively small size of far-IR detectors compared to their visible and near-IR equivalents, it could map large areas of sky efficiently to faint limits. The telescope was designed to give diffraction-limited images - resolution 6 arcseconds - at 90 microns but, in space, it actually performs significantly better than this, with diffraction-limited images being seen as short as 70 microns, with a FWHM of 5.5 arcseconds at this wavelength. In mapping mode, at fast scan speeds there is, as is logical, some degradation of the PSF in the direction of scan.

3.1. General aspects

The Herschel Space Observatory covered the wavelength range from approximately 55 - 672 microns. This corresponds to the maximum of emission for black bodies in the range from 5-50K approximately. Hence Herschel was be best suited to observing icy outer solar system objects and cool and cold dust in the universe, both in the rest frame and redshifted. A prime objective has been to study the formation of galaxies in the early universe, as cool dust is an excellent tracer of star formation. The Herschel range was also the one at which cool and cold gases emit their strongest lines, meaning that Herschel was also a superb laboratory for examining the chemistry of planetary atmospheres and of the interstellar medium.

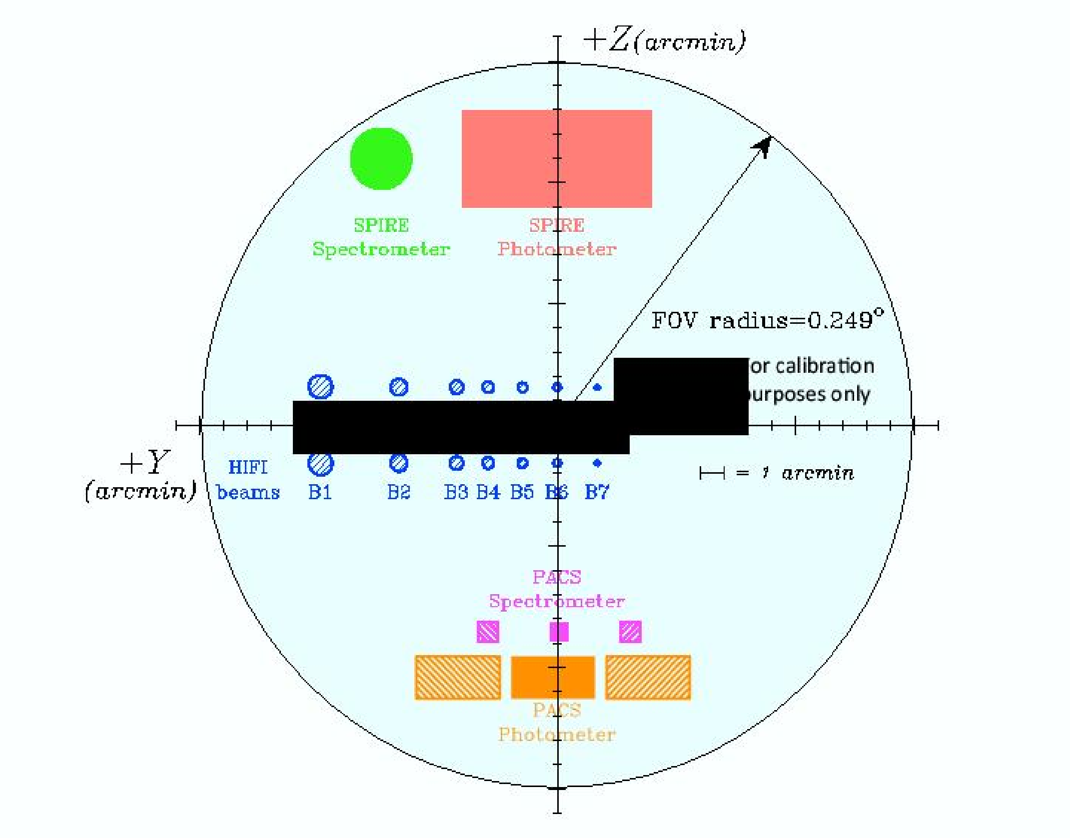

Figure 3.1. The Herschel Focal Plane showing the field of view of the telescope and the positions of the different instrument apertures within it. The purpose of the telescope SIAM (Spacecraft Instrument Alignment Matrix) was to provide exact offsets to each of the apertures from the reference position used in pointing so that the target is accurately centred in the aperture of interest.

The Herschel Focal Plane is shown in Figure 3.1. The positions of the different instrument arrays and apertures are labelled. The full, unvignetted field of view is approximately half a degree. The purpose of the spacecraft SIAM (Spacecraft Instrument Alignment Matrix) was to set the correct off-sets to allow Herschel to point accurately to each individual aperture on demand from the target reference position. Accurate determination of the SIAM was a critical part of commissioning the telescope and its refinement through the mission was important to take advantage of the improved seeing performance.

3.2. Photometry with Herschel

3.2.1. PACS Photometer science

3.2.1.1. Instrument capabilities

The full wavelength range of Herschel was covered by six broadband ($\Delta\lambda/\lambda$=3) filters. In SPIRE, all three filters (250, 350 and 500 µm) were imaged simultaneously on three spiderweb bolometer arrays. PACS users were able to image with a "red" (130-210 µm) and a "blue" (either 60-85 or 85-130 µm) filter simultaneously on two bolometer arrays.Colloquially, but inaccurately, the 130-210, 85-130 and 60-85 micron filters are often referred to as the PACS "red", "green" and "blue" channels.

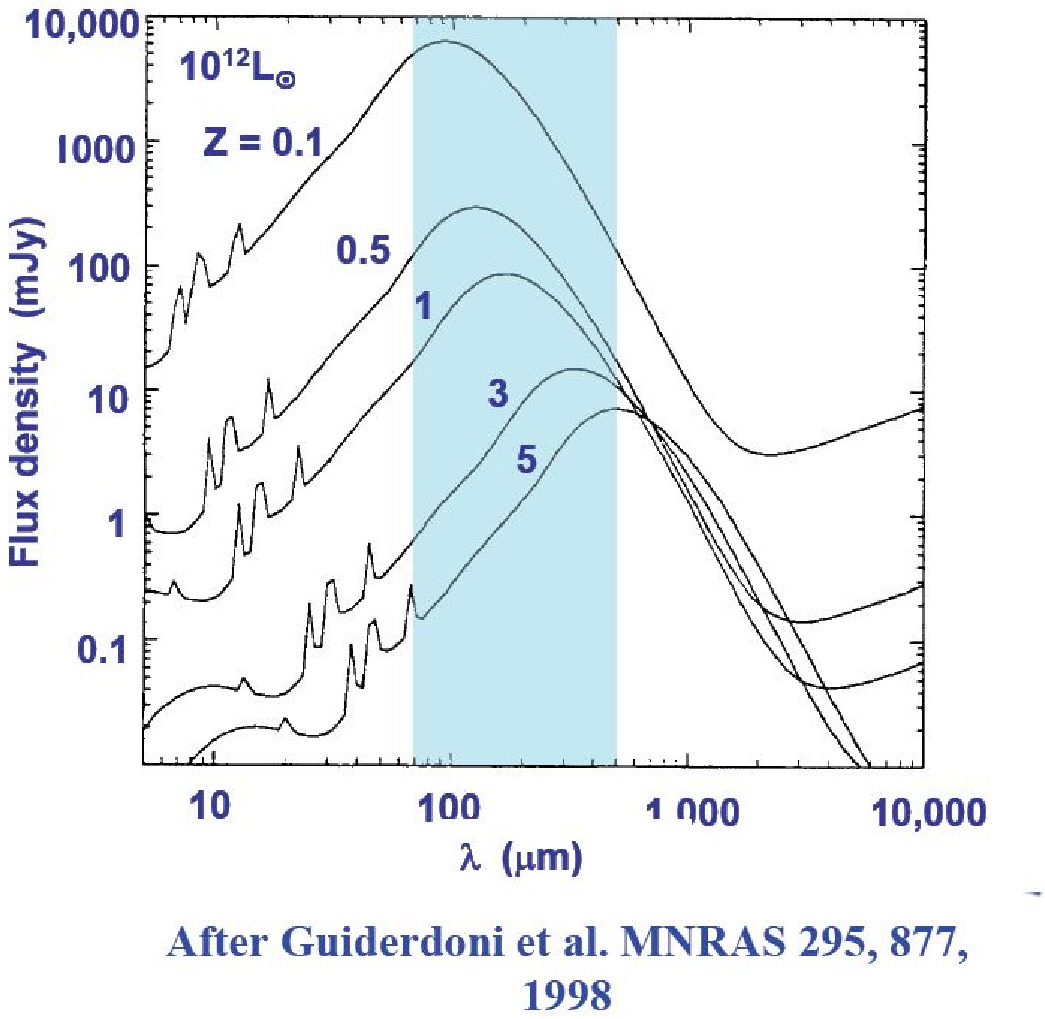

Figure 3.2. The black body spectrum of the far infrared dust peak of a star-forming galaxy of 10 solar masses showing the effect of redshift on the spectrum. At z=5 the peak moves to the long-wavelength end of Herschel coverage with SPIRE. The result of this is that the colour of a galaxy in a SPIRE RGB image

gives a good estimate of its redshift and thus distance, with strongly red galaxies in SPIRE image being the most distant.

This has made Herschel a superb instrument for multicolour surveys. SPIRE could image a square degree of sky to an instrument noise equal to the extragalactic confusion limit (1 sigma) in 2.3 hours: a 60x60 arcminute large map with SPIRE with 2 repetitions would reach the confusion noise at 1 sigma in this time and, of course, much larger areas of the sky to a lesser sensitivity. Due to the overheads involved in scan maps (i) making large maps was more efficient than making small maps and (ii) making square maps, or maps with a small aspect ratio was somewhat more efficient than making long, thin strip maps, as the overheads implicit in long scans were reduced: a one square degree map made as a 60x60 arcminutes square would take 8186s, including overhead. A one square degree map taken as a 200x18 arcminutes strip would take 10564s, including overhead.

The main imaging capabilities are summarised in Table 3.1. As dust is a strong tracer of star formation (see Figure 3.2), one of Herschel's greatest strengths has been the possibility to study the history of star formation in the universe. By combining PACS and SPIRE data, users are able to follow the dust emission signature of starbursts redshifted to increasing wavelengths in ever more distant galaxies; this has made Herschel an enormously powerful facility for studying the formation and evolution of galaxies.

Observations with Herschel give a new insight into the process of star and planet formation. Herschel could study both the processes of star formation in molecular clouds and the debris disks that are the tracer of planetary system formation in young stars. Prior to Herschel's launch, few debris disks were known; observations with Herschel, with its wide wavelength coverage, have allowed many more to be detected and studied (Figure 3.3). Similarly, Herschel observations are valuable in the study of the later phases of stellar evolution, particularly circumstellar shells (Figure 3.4), mass-loss in general and stellar winds.

Finally, Herschel is a powerful tool for studying the physics of the more distant and colder objects of the solar system: such as the atmospheres of the giant planets (Figure 3.6), their icy satellites, cometary nuclei and cometary atmospheres. Herschel observations are permitting the albedos and thus the surface conditions and diameters of these bodies to be measured with great precision.

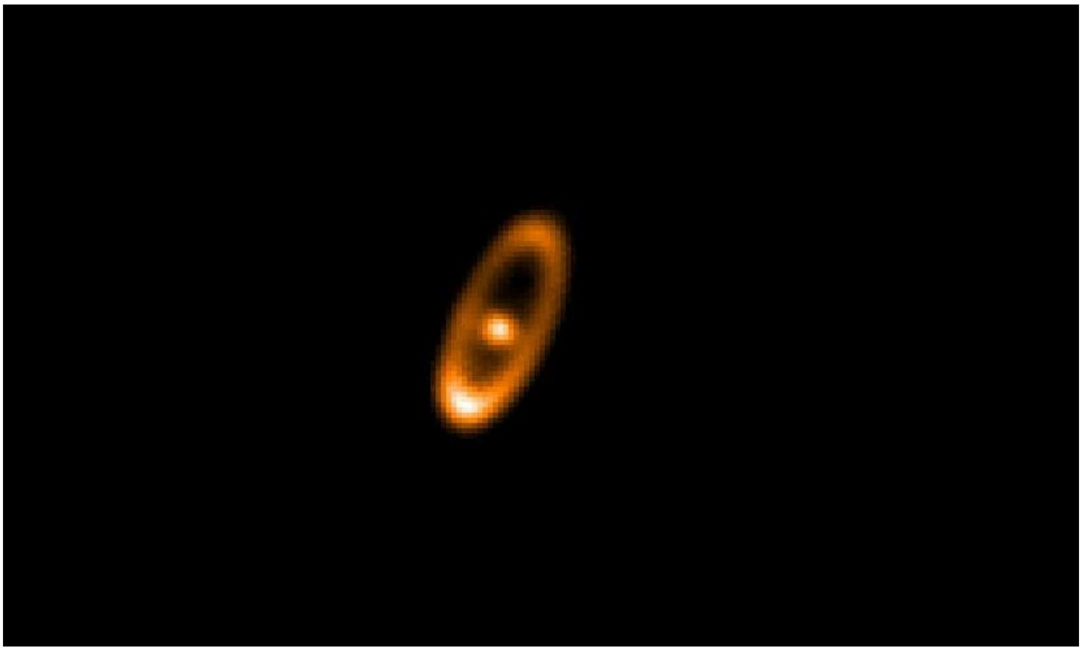

Figure 3.3. A PACS 70 micron image of the debris disk around Fomalhaut. The debris disk is well resolved. This is one of the few cases where Herschel also detects a hot star as, at 16 light years distance, Fomalhaut is close enough to be strongly detected at short wavelengths.

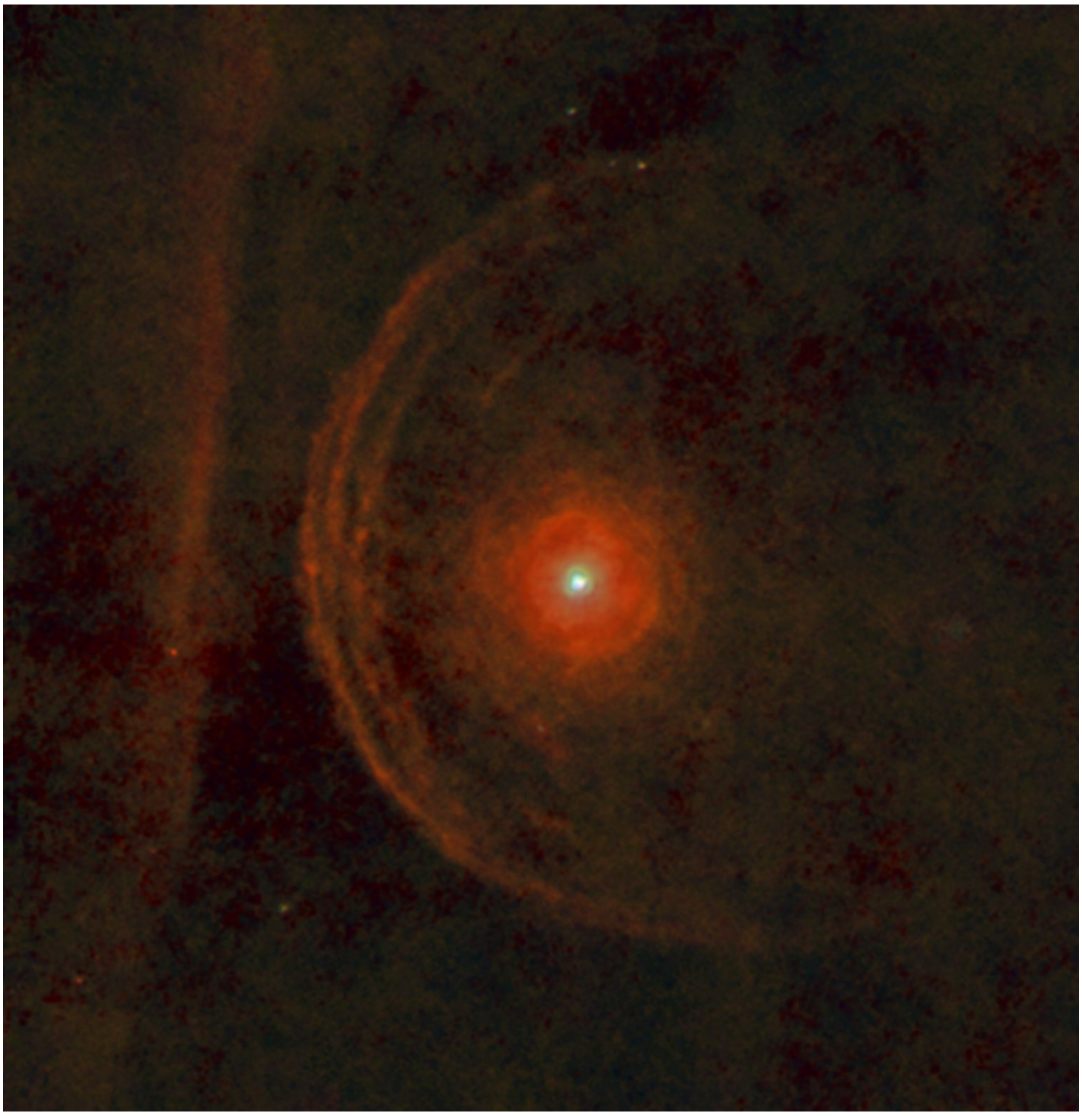

Figure 3.4. An RGB image of Betelgeuse and its environment constructed from PACS images at 70, 100 and 160 microns. North is to the top left and east to the bottom left in this view. A series of arcs can be see to the left (northeast) of the star; these are shells of material shed by the star in its supergiant phase. An intriguing feature of this image is the vertical bar to the left, which it is hypothesised may be the edge of an interstellar cloud that is being illuminated by the star; if so, the outer shell will collide with it in approximately 5000 years.

Table 3.1. The main imaging capabilities of PACS and SPIRE. Please note that the wavelength range of detector sensitivity is approximate and the instrument sensitivities depend on the observing mode, so the values given are only orientative: please consult the relevant observing manual for more detailed values.

	PACS	SPIRE
Wavelength range	60-210 µm	200-670 µm
Field of view	1.75x3.5'	4x8' (unfilled)
Pixel size	3".2 (60-130 µm), 6".4 (130-210 µm)	18".2 (250 µm), 24".9 (350 µm), 36".3 (500 µm)
Typical sensitivity (5σ/1hr)	5 mJy (70/110 µm bands), 10 mJy (160 µm band)	Will reach the confusion limit (see below)
Confusion limit (ideal case)	<0.1 mJy (70 µm), 0.27 mJy (100 µm), 0.92 mJy (160 µm)	5.8mJy (250 µm), 6.3mJy (350 µm), 6.8mJy (500 µm)

Filters	60-85 or 85-130 µm and 130-210 µm (simultaneous)	250, 350 and 500 µm (simultaneous)

3.2.1.2. PACS red matrix issues from OD-1375

One matrix of the PACS 160 micron channel was found to be unresponsive in data taken on OD-1375 (February 16th 2013). Despite considerable efforts by the PACS ICC, it could not be recovered, leading to the loss of 50% of the detector in this channel. This affected the spatial coverage offered by the channel, requiring mapping observations to be re-designed to allow them to be observed efficiently. The consequence was a small reduction in spatial coverage of 160 micron maps compared to 70 and 100 micron maps, plus a loss of approximately 30% in sensitivity in this channel from then until the end of the mission. Observations in Parallel Mode (Section 3.2.2) taken after the failure of the matrix show uneven coverage and severe striping in the 160 micron band due to the lack of the coverage from the second matrix.

3.2.2. Using SPIRE and PACS in parallel

Herschel offered a Parallel Mode for users who wished to carry out large-scale mapping programmes with a wide range of wavelength coverage.

3.2.2.1. The benefits of using Parallel Mode

Figure 3.5. An example of use of the SPIRE PACS Parallel Mode. This is a Galactic Plane image from L=316 degrees, just north of Alpha Centauri, taken by the Hi-GAL project, as part of their 360 degree survey around the Galactic Plane. The use of bands from 70-500 microns in a single observation allows the different warm and cool dust emission to be identified, sampling a wide range of temperature and of physical conditions. The ability of Parallel Mode to cover large areas of sky rapidly was vital to the success of the Hi-GAL mapping.

Parallel Mode allowed observers to use both SPIRE and PACS simultaneously in a fast (60 arcsec/s) and nominal speed (20 arcsec/s) scanning mode, to cover very large areas of sky quickly in all three SPIRE bands and in two of the three PACS bands, to a modest sensitivity. This mode was intended to make ambitious, multi-band, large area mapping programmes more efficient than carrying them out individually with each instrument in turn (see Figure 3.5). In this mode SPIRE is the prime instrument and thus the driver in defining observations and PACS data, which is much shallower, particularly in the 70 micron band, should be treated more as a "bonus" to observers.

3.2.2.2. The limitations of using Parallel Mode

SPIRE and PACS pointed at different places on the sky separated by 21 arcminutes. This meant that this mode was extremely inefficient at mapping small areas of sky. Although a minimum area of 30x30 arcminutes for a Parallel Mode map was permitted by HSpot, Parallel Mode was not recommended for observing any area of sky smaller than one square degree and only became efficient for mapping for even larger areas than this. In contrast, observing with SPIRE alone was highly efficient for mapping large areas of sky quickly such that a number of Parallel Mode programmes moved to SPIRE-only mapping instead.

While the SPIRE integration gets to a depth which is fairly close to the confusion limit, the depth of exposure is relatively less for PACS and the flux limit attained in PACS is quite shallow and thus suitable only for bright sources and regions, hence Parallel Mode photometry should not be regarded as an adequate substitute for dedicated PACS scan maps in any area of the sky.

Although very large areas of sky could be mapped quickly at the high scan speed, the speed of scan was sufficiently high that some telescope movement occurred before detector readout was complete, giving rise to a small degree of PSF smearing in the scan direction.

Parallel Mode entailed having both PACS and SPIRE cool simultaneously and is thus was more demanding in terms of helium usage. As a rule of thumb, for each three parallel cooler re-cycles (6 days of Parallel Mode use) we reduced mission lifetime by one day with respect to using just PACS or SPIRE alone. For efficiency of helium consumption, using Parallel Mode was thus discouraged unless the increased wavelength coverage was genuinely required scientifically and made practical sense.

3.3. Spectroscopy with Herschel

Herschel offered two types of spectroscopic capability. PACS and SPIRE offered low to intermediate resolution spectroscopy covering the full Herschel wavelength range. HIFI offered high-resolution spectroscopy over the range from 157-625 µm (480-1910 GHz) using heterodyne techniques, although there is a small gap in coverage from 213-240 microns (1272-1430 GHz), between the HIFI 5b and 6a sub-bands: Figure 3.7 shows this gap in coverage, but also illustrates the extraordinary capabilities of HIFI for detecting spectral lines. Users were thus be able to select a wide range of resolutions from $\Delta\lambda/\lambda$=20 to $\Delta\lambda/\lambda$=10 000 000 according to the brightness of their source and the science that was required. The main spectroscopic capabilities are summarised in Table 3.2.

The short wavelength cut-off for Herschel is a matter of definition. For PACS the detector sensitivity below 55 microns was too low to be of practical use in most cases and so this value is given as a limit here. However, there is an O III line at 52 microns that is strong enough to observe in some targets. The sensitivity of PACS in the 50-55 micron range was updated in HSpot in Autumn 2011, as the initial values given were found to be too optimistic by a factor of about 3; this correction though made systematic observation at wavelengths shorter than 55 microns even more impractical than originally envisaged.

In its highest resolution mode Herschel offered a velocity resolution as high as 0.3km/s. The wavelength range covered by Herschel has many thousands of lines of water, atomic transitions and organic molecules, as shown by the extraordinary spectrum in Figure 3.7, observed by the KPGT_ebergin_1 Key Programme, in which 75 000 lines have been detected so far, although only some 20 000 have been identified to date. This allowed Herschel to study the chemistry of the interstellar medium, tracing water and organic molecules in molecular clouds. Herschel was also be able to study the chemistry of solar system bodies such as the atmosphere of Mars and the comas of comets in unprecedented detail.

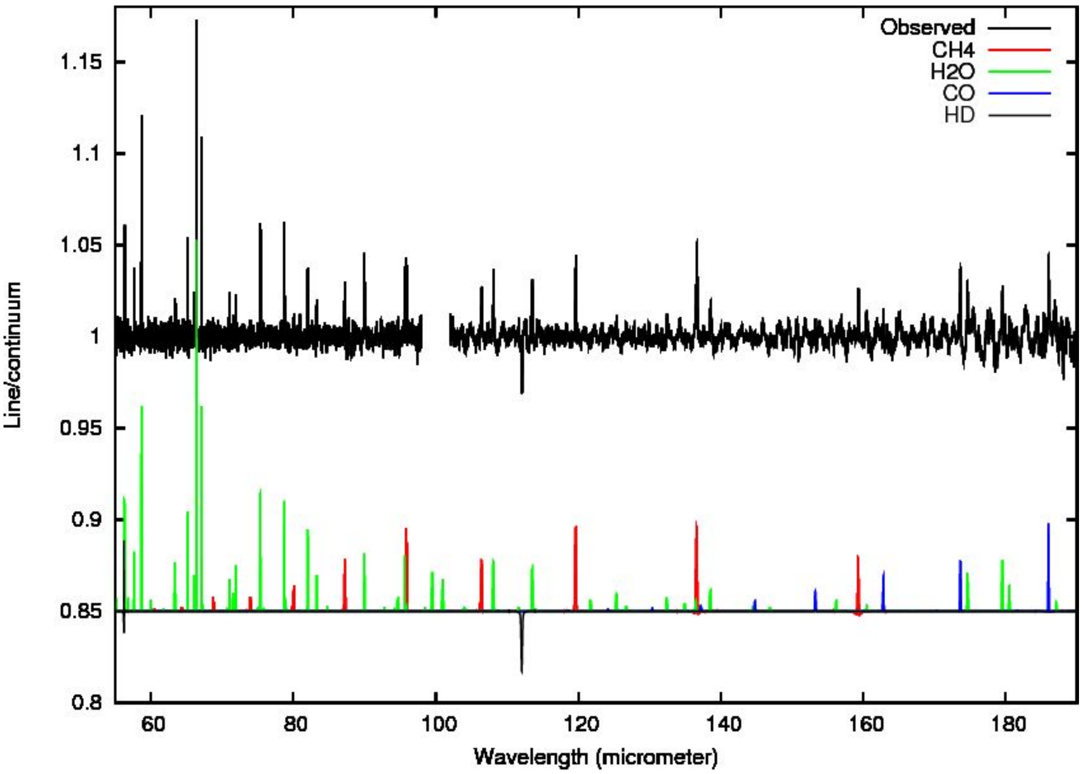

Figure 3.6. A PACS composite spectrum of Neptune from 57 to 190 microns (upper pane) and the corresponding line identifications (lower pane). The spectrum is dominated by water emissions at short wavelengths, with methane, carbon moxide and deuterated hydrogen important above 90 microns. From [Lellouch et al.: 2010, A&A, 518, L152]

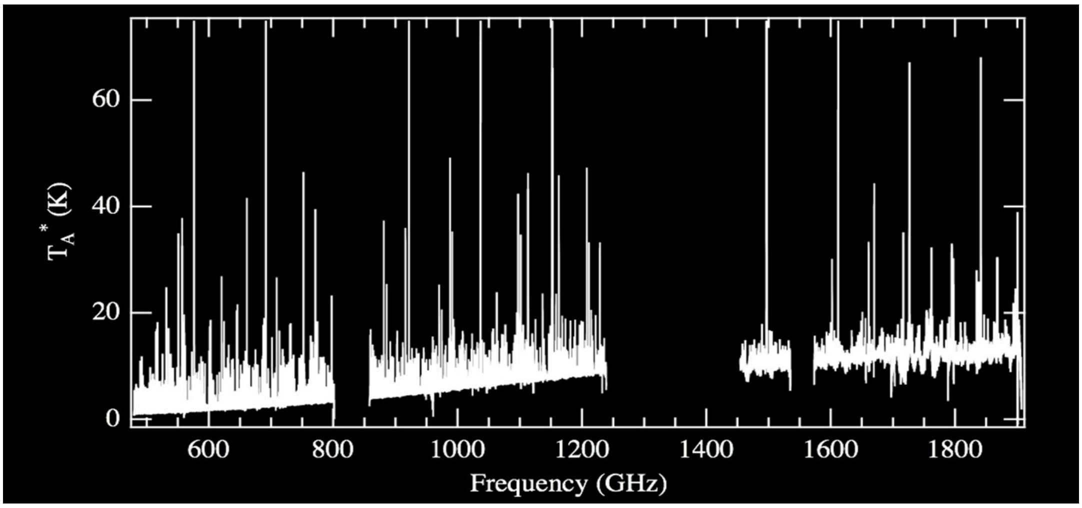

Figure 3.7. A full spectrum scan of Orion KL taken with HIFI for KPGT_ebergin_1. This spectrum shows the extraordinary number of spectral lines visible in this source. A total of 75 000 lines have been detected in the spectrum, only 20 000 of which have been identified so far, of which 5000 are from methanol. The gap in the spectral coverage between bands 5b and 6a is clear in this view with its linear frequency scale.

All three instruments had a mapping capability in spectroscopic mode, even though HIFI's was somewhat limited, although by no means negated, by the fact that its detector had only a single pixel. PACS and HIFI could scan the detectors across the sky, accumulating spectroscopic data along the length of the scan. SPIRE could not do that, instead it used a beam-steering mirror to make filled maps. All three instruments could make a raster map in spectroscopic mode. This allowed a spectroscopic survey to be made either of a region that has been mapped in imaging mode, such as a cluster of galaxies, or across a known extended source such as a molecular cloud.

Table 3.2. The main spectroscopic capabilities of PACS, SPIRE and HIFI. For more details please check the relevant instrument manual.

	PACS	SPIRE	HIFI
Wavelength range	55-210 μm	194-313 and 303-671 μm	157-213 and 240-625 μm (with gap)
Field of view	47x47"	2.0' (unvignetted)	Single pixel (see below)
Pixel size	9"	17", 29" (varies across the bands)	39" (488GHz), 13" (1408GHz)
Sensitivity (5σ/1hr, point source)	$2x10^{-18}$ Wm^{-2} (130 μm, 1st order), $5x10^{-18}$ (70 μm, 3rd order). Continuum: 100 mJy (1st order), 250 mJy (3rd order)	$1.0\text{-}2.2x10^{-17}$ Wm^{-2} (high resolution), 40-88mJy (low resolution) [5 sigma/ 1 hr]	"A few" mK (Band 1a) to 100mK (Band 7b), 1 sigma/1hr
Resolution	900-2100 (1st order, 102-210 μm), 1800-3000 (2nd order, 72-98 μm), 2600-5400 (3rd order, 55-72 μm)	20-1000	$1000\text{-}10^{7}$

Note

For the latest information on instrument sensitivities please check the Herschel website at http://herschel.esac.esa.int/.

Note also that the PACS sensitivity below 57 microns is very low, although HSpot permitted the entry of line observations at wavelengths down to 51 microns. It is a matter of definition where the short wavelength cutoff of PACS is set. Some successful observations of the strong O III line at 51.81 microns were carried out, but only on the brightest available targets.

3.4. Obtaining observing time with Herschel

3.4.1. The standard method for obtaining observing time with Herschel

The standard method of obtaining observing time with Herschel was from the regular Announcements of Opportunity and their associated Calls for Proposals, the last of which was the OT2 Call that closed in summer 2011. However, for urgent observing requests that could not reasonably be covered by the Calls for Proposals, or that could not have been anticipated when a Call for Proposals was open, a quantity of observing time was reserved for Director's Discretionary Time (DDT). This could be obtained, when the circumstances justified it, as a DDT proposal (non urgent) or, for urgent observations, a Target of Opportunity (ToO) request.

There were three Announcements of Opportunity, each consisting of a Guaranteed Time and an Open Time Call for Proposals. Approximately one third of all Herschel time was reserved as Guaranteed Time, mainly for the instrument consortia who built the instruments, with a smaller amount distributed among the Herschel Science Team members.

3.4.1.1. The Key Programme Call

It was recognised pre-launch that it was essential to dedicate a large fraction of observing time early in science operations to a small number of core science programmes proposed by large consortia, which would serve to populate the Herschel observations database with large numbers of targets that could be scheduled quickly, as soon as observing modes were released. These programmes would serve as training for a substantial number of people in the community in proposal preparation under the guidance of scientists familiar with the Herschel instrumentation and its capabilities. To be classified as a Key Programme, a proposal needed to be for a minimum of 100 hours of observing time.

The Announcement of Opportunity was made on 1st February 2007, with a deadline of 4th April 2007 for Guaranteed Time Key Programme (KPGT) proposals. However, due to a total power failure both of primary and emergency power supply at ESAC on the morning of closure that took down all computer services across the site, it was decided to extend the deadline for 24h to 5th April 2007. The call for Open Time Key Programme (KPOT) proposals opened on July 5th and closed on 25th October 2007.

A total of 11258h of observing time were awarded, although in the end only 10826h were actually used (this difference is due to changes in the time estimator making observing more efficient and modifications to programmes based on in-flight experience, leading to Guaranteed Time originally assigned to KPGT proposals being transferred to other, later programmes). The list of approved programmes, 21 Guaranteed Time and 21 Open Time, is shown in:

http://herschel.esac.esa.int/Key_Programmes.shtml.

The sky distribution of Key Programme AORs is shown in Figure 3.8, showing the remarkably homogeneous coverage of the sky by the 11650 Key Programme AORs approved by HOTAC and entered at the end of Phase 2 of proposal entry, easing the potential problems of inhomogeneous availability of targets through the year making scheduling inefficient.

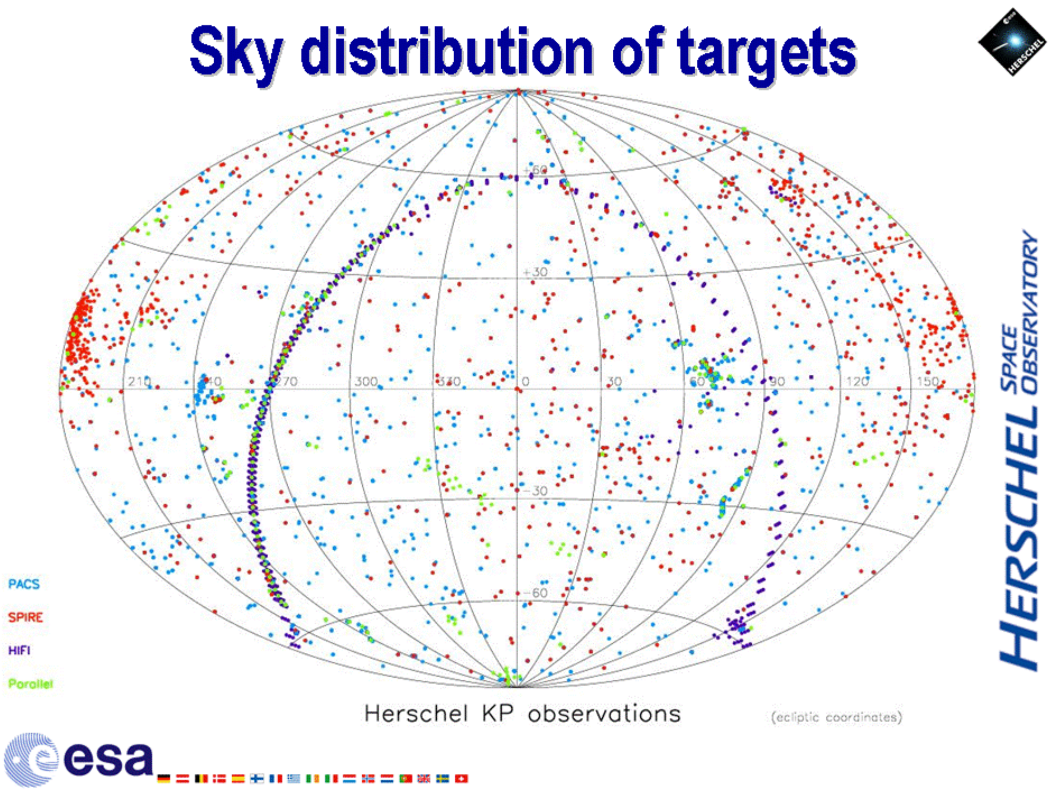

Figure 3.8. The sky distribution of AOR centres for the 11650 accepted Key Programme AORs. Although there is a concentration both in the Galactic Plane and at the Galactic Poles, the overall distribution is remarkably homogeneous.

The time awarded by HOTAC was 57% of the nominal Herschel science time available for the mission of 19776h, consisting of 93% of the nominal Guaranteed Time and 40% of the nominal Open Time for the mission.

Given that the Key Programme Call closed more than a year and a half before launch, the AORs were based only on a theoretical knowledge of their likely behaviour in space, obtained from laboratory testing. In reality it was accepted that the time estimates were likely to be within 10% only of the final in-flight values, with the possibility that times could either increase or decrease in flight as the observing modes were tested and refined, thus there was the *potential* for as much as 1000h change between the original HOTAC awards and the actual time required to execute the observations.

Some important changes in the way that observations were executed were made very quickly. For example, it was found early in the mission that the slew time estimator was extremely conservative and that around 50s could be saved on all slews. As scan maps were made as line scans, with large numbers of slews, this led to a substantial saving in overheads for scan maps and a considerable saving in the time required to execute a map when implemented in November 2009. Similarly, it was found that the settling time for individual pointings in raster maps was overly conservative, with a minimum 8s settling time and larger values for bigger raster steps applied during Key Programme submission; this was reduced to a flat 5s (and, later, 3s for slews between line scans), even before launch, with substantial savings to many programmes. In contrast, it was found necessary to modify the way that PACS spectroscopy AORs were taken, leading to a significant increase in observing time for some proposals. Similarly, changes in the HIFI sequencer could either increase or decrease observing time according to the sub-mode and the band required. And the change of PACS photometry from point-source photometry mode to mini-scanmaps made a fundamental change to the execution of a large fraction of all the approved AORs.

3.4.1.2. The First Open Time Call

The first in-flight Announcement of Opportunity (AO1) was made on February 25th 2010. This distributed most of the remaining Guaranteed Time (555h), plus 6577h of Open Time, of which 4988h were Priority 1 ("guaranteed" execution) and 1589h were Priority 2 time (fillers). The Guaranteed Time (GT1) Call opened on February 25th 2010, with it submission deadline on March 31st 2010. The Open Time Call (OT1) opened on May 20th 2010, with a submission deadline of 22nd July 2010.

A total of 33 GT1 and 241 OT1 proposals were accepted for execution, of which 176 of the accepted OT1 proposals were awarded Priority 1 time. The list of approved programmes is shown in http://herschel.esac.esa.int/AO-1_Programmes.shtml.

As a matter of policy it was decided that, as fillers, no constraints could be accepted on Priority 2 observations. This meant that no observations with timing constraints could be accepted as Priority 2 (a timing constraint was a scheduling window, an orientation constraint, a follow-on, or any other kind of scheduling restriction other than a concatenation). Solar system targets were defined to be time-constrained by nature. Proposals with time-constraints had either to be ranked high enough to obtain Priority 1 time, or be rejected.

3.4.1.3. The Second Open Time Call

The Second Announcement of Opportunity (AO2) was made on April 7th 2011. This distributed the very small amount of remaining Guaranteed Time (362h) plus 7590h of Open Time, of which 3420h were Priority 1 ("guaranteed" execution) and 4170h were Priority 2 time (fillers, with an expected execution rate of around 18% - in fact, the final degree of completion of Priority 2 time was 48%). The Guaranteed Time (GT2) Call opened on April 7th 2011, with it submission deadline on May 12th 2011. The Open Time Call (OT2) opened on June 9th 2011, with a submission deadline of 15th September 2011.

A total of 32 GT2 and 373 OT2 proposals were accepted for execution, of which 181 of the accepted OT2 proposals were awarded Priority 1 time. The list of approved programmes is shown in http://herschel.esac.esa.int/AO-2_Programmes.shtml.

The Proposal Handling System (PHS) received a rapidly increasing load during the last few days before closure. Figure 3.9 shows the evolution of load during the last 30 hours before closure. The load peaked at one proposal submission every 15 seconds shortly before closure. Most of the load was due to re-submissions of already submitted proposals as users followed advice to make an initial security copy submission and then fine-tune it, re-submitting a new version after each revision.

However, during the last hour before closure, newly submitted proposals were also arriving at a rate of approximately one every two minutes.

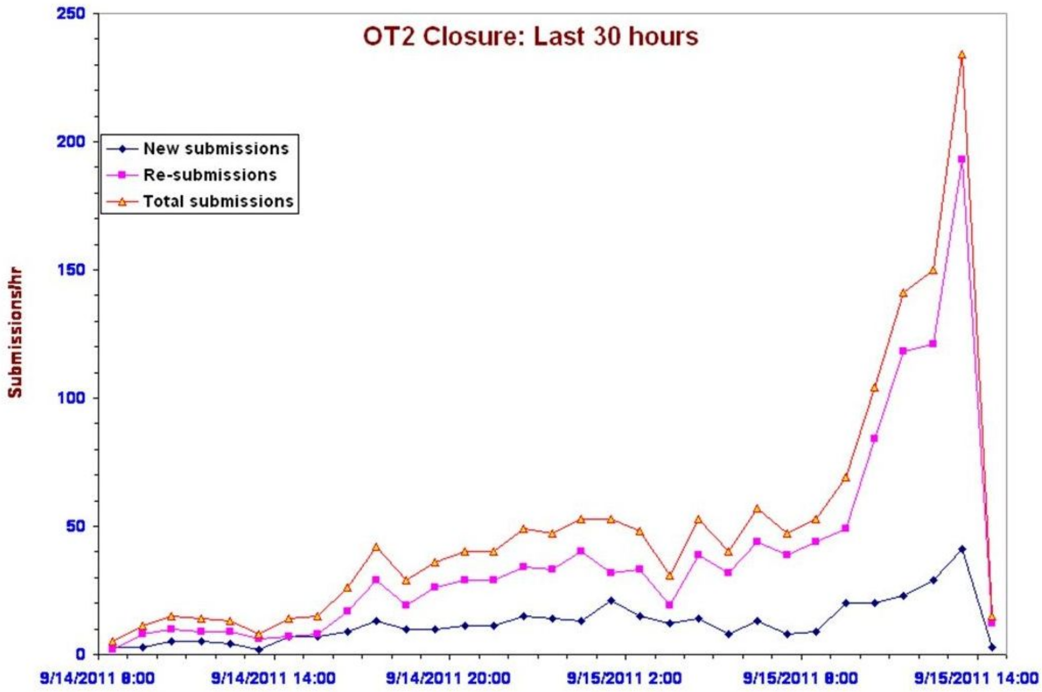

Figure 3.9. Proposal Handling System load at closure of the OT2 Call. Shortly before closure the system was receiving and processing a proposal submission every 15 seconds.

Comparison of system load at closure of the OT1 and OT2 Calls showed that they were remarkably similar (see Figure 3.10) shows the evolution of load during the period that submissions were being received. The pattern observed during the last two weeks of the Call is almost identical. The slight discrepancies in the load between the two Calls prior to this can be explained by the fact that the OT1 Call closed before the August holiday period and the OT2 Call closed after it.

Save for cases where a target had visibility issues, all OT1 and OT2 Priority 1 observations had to be completed by December 31st 2012 to ensure guaranteed execution by EoH. In reality, efficient scheduling meant that OT1 and OT2 Priority 1 scheduling was essentially complete by early December 2012, allowing bulk scheduling of Priority 2 observations to start early.

3.4.1.4. The "Must Do" Call

As EoH approached, the Herschel Mission Scientists undertook the task to review the approved to-be-executed observing programmes with the objective of identifying potential "gaps" that could and should be filled. The aim was to identify any science that Herschel must do to avoid leaving gaps in its legacy. After a careful review, only a small number of such "Must-Do" observations were identified, showing that the Herschel legacy was, to a very large degree, complete, so the process was opened to the general community. The announcement was made on 25 May 2012, with a 14 June 2012 deadline.

Only completely new programmes that had never been proposed to HOTAC and that were deemed to be of very high scientific priority were accepted for execution as Must Do programmes. A total of 107.6h from seven proposals was approved after technical assessment and duplication checking; 43% of the approved time was awarded to a single, Planck follow-up programme. The list of approved programmes is shown in http://herschel.esac.esa.int/Must-Do_Programmes.shtml.

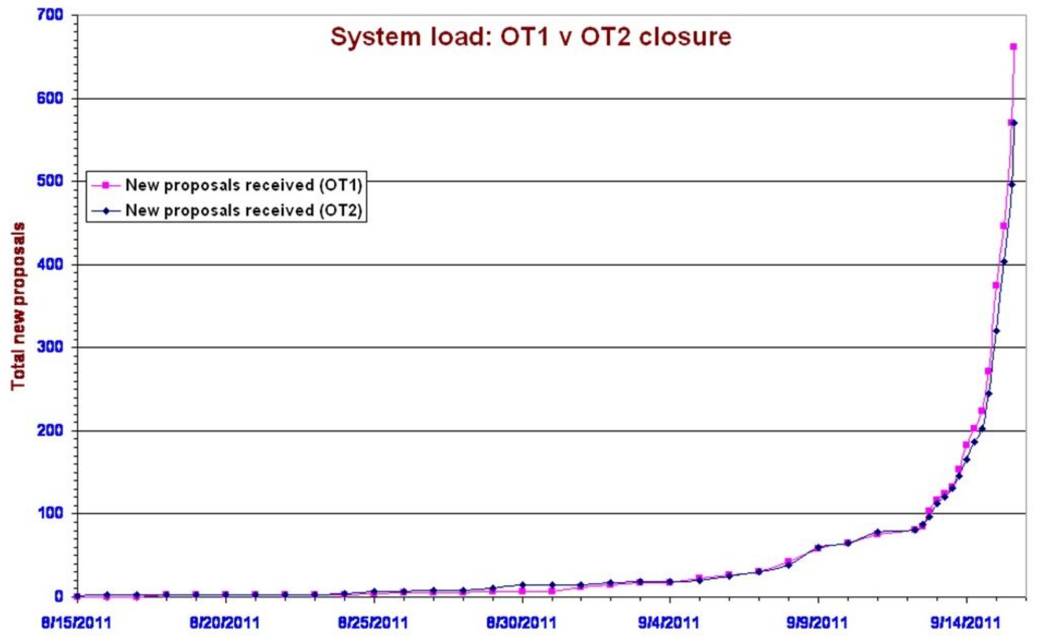

Figure 3.10. Comparison of the Proposal Handling System load at closure of the OT1 and OT2 Calls. The pattern is almost identical, demonstrating that the Herschel community are creatures of habit!

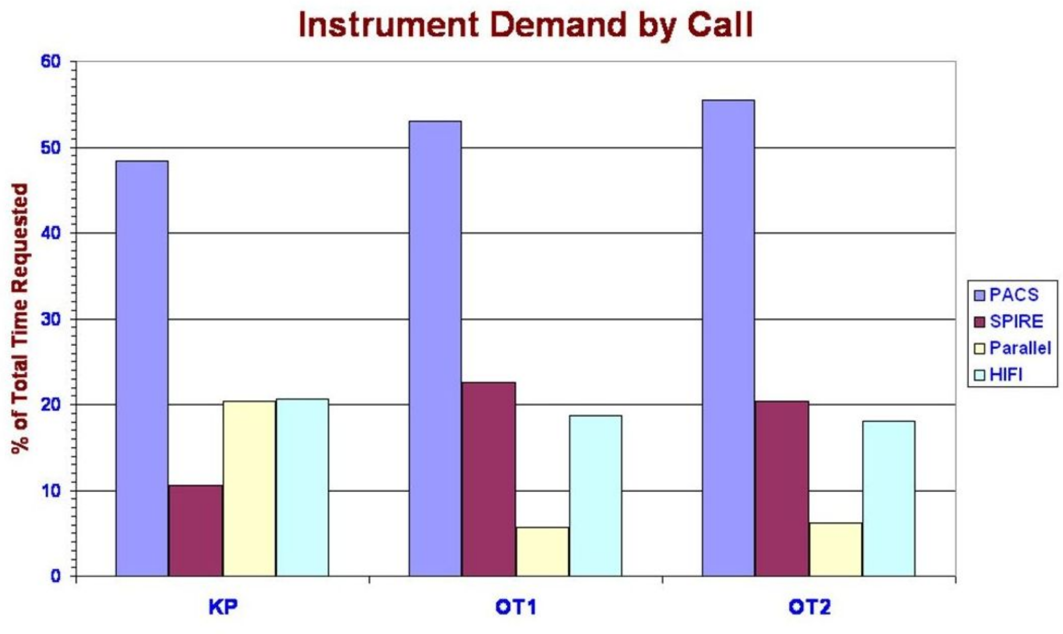

Figure 3.11. Comparison of the fraction of time requested for each of the Herschel instruments (counting SPIRE PACS Parallel Mode as a separate instrument) as a function of the three major Calls for Proposals during the mission. The overall level of demand for the instruments changes very little from Call to Call although it is noticeable that demand for Parallel Mode decreased sharply after the Key Programme Call as it became evident that, for many programmes, it was more efficient to use SPIRE only. In flight knowledge of the instrument performance showed that Parallel Mode had only very limited applications for cosmological programmes due to the far shallower depth reached by the PACS component of the data.

The first Must Do observations were scheduled on OD-1236. Scheduling of Must Do was essentially complete at EoH, with only a tiny number of observations with visibility beyond EoH that were not executed.

3.4.1.5. The Third Guaranteed Time Call

With Boil-Off well delayed with respect to predictions, there was testimonial announcement of a third Guaranteed Time Call at the end of April 2013, with no fixed deadline. The aim was to top-up the database with additional observations as they were submitted and accepted to improve scheduling efficiency, which would be adversely affected if helium continued through towards the end of Spring 2013. As boil-off happened within two days of the announcement of the GT3 Call, although three proposals were received before boil-off made the Call redundant, no GT3 observations were scheduled.

3.4.1.6. Evolution of instrument demand during the mission

It was expected that instrument demand might change through the mission as the community grew to know the instrumentation better and passed from survey programmes to follow-up. In fact though, the demand for instruments proved to be remarkably stable through the mission. Figure 3.11) shows the evolution of the demand for each intrument -- treating SPIRE and Parallel Mode as separate entities which, in practice, they were not -- while Figure 3.12) shows the equivalent plot for the different sub-modes. The biggest single change is the shift from Parallel Mode to SPIRE photometry as the performance of Parallel Mode and its issues (see Section 3.2.2) became better understood.

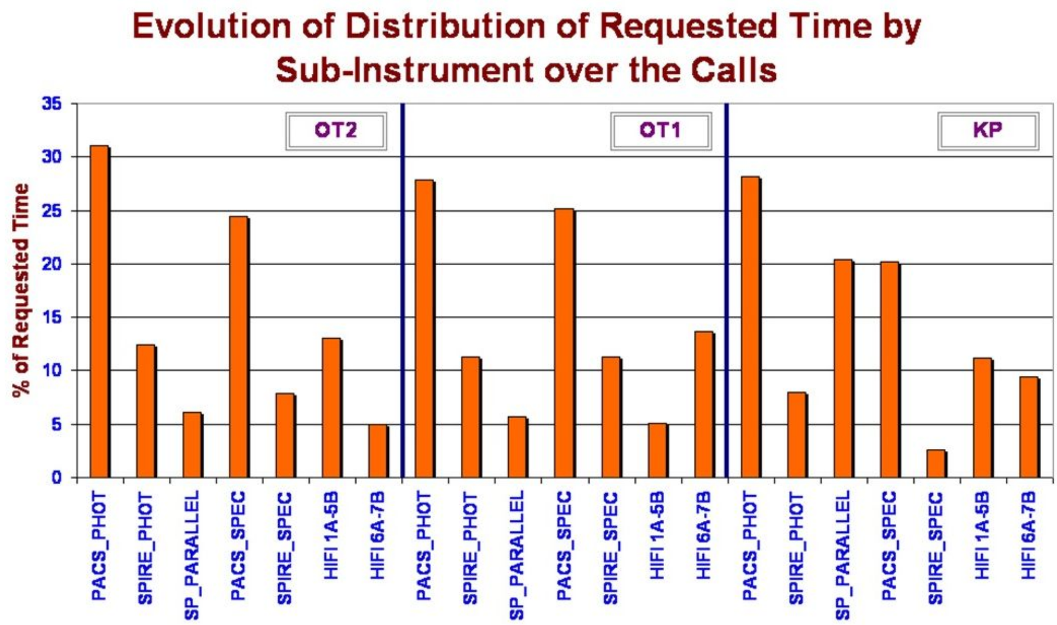

Figure 3.12. Comparison of the fraction of time requested for each of the Herschel sub-instruments as a function of the three major Calls for Proposals during the mission. The relative demand for the different sub-instruments was remarkably stable between the closure of the initial Call in 2007 and the final Call in 2011.

3.4.2. Urgent scheduling requests, DDT proposals and ToOs

The standard method of obtaining observing time with Herschel was from the regular Announcements of Opportunity and their associated Calls for Proposals, the last of which was the OT2 Call that closed in summer 2011. However, for urgent observing requests that could not reasonably be covered by the Calls for Proposals, or that could not have been anticipated when a Call for Proposals was open, a quantity of observing time was reserved for Director's Discretionary Time (DDT). This could be obtained, when the circumstances justified it, as a DDT proposal (non urgent) or, for urgent observations, a Target of Opportunity (ToO) request.

Although 5% of all Herschel nominal science time (1000h) was "reserved" for DDT proposals, only 1.5% of the final Herschel observing time was awarded in the form of Director's Discretionary Time.

3.4.3. Ground station access to Herschel

The Herschel Space Observatory was fundamentally an "off-line" mission and was designed as such, so reacting rapidly was problematic. "Off-line" means that, instead of being in permanent contact with Herschel there was a Daily TeleCommunications Period (DTCP), normally for 3 hours each day, when the data stored on board had to be downlinked and new observations uplinked. The DTCP was within a few hours of local midnight at the ground station (prime was New Norcia, back-up was Cebreros), as Herschel was always close to the Opposition Point in the sky and thus at optimal altitude to be contacted for a few hours either side of local midnight from any ground station. The DTCP time varied according to whether Planck or Herschel was visible to the antenna first - every few months their orbits crossed over and the lead satellite in the DTCP changed - and also to demands on Ground Station resources from other missions that sometimes forced the communications window to be shifted further from the optimum time, based on a strict hierachy of Operational priorities for mission ground station access.

The off-line nature of Herschel meant that everything had to be ready for uplink before the DTCP started. This DTCP start time defined an unbreachable barrier and everything was calculated backwards from this moment.

At each DTCP, observations are uplinked for execution from 27 to 51h ahead. This means that even if a Ground Station pass were missed due to a communication problem, there were always enough observations in the on-board memory for Herschel to continue operating until the end of the DTCP period of the following day. All processing time for the planning file of observations to prepare it for uplink had be added to this minimum 51h. In practice, only two DTCPs were missed during the entire mission, one due to a sudden, very heavy snowfall at the Cebreros ground station that left the antenna unusable and the other due to a technical issue at New Norcia.

Ground station access to Herschel was thus the fundamental limitation to the Herschel reaction time to urgent scheduling requests. Assuming that the science was important enough to trigger a re-planning, the reaction to an urgent scheduling request was totally dependent on whether or not the re-planned observations could be processed and ready in time for the start of the target ground station pass. If the answer is not clearly "yes", it was unlikely that the request would be accepted.

3.4.4. Herschel reaction time to urgent scheduling requests

There were a series of factors that decided how quickly Herschel could react to an urgent scheduling request. These are described briefly below. The faster the reaction required, the bigger the effort that was required to achieve it and the greater the knock-on effects and risk to normal spacecraft operations that it entailed. In general, any change to the observing schedule made less than 3 weeks before execution of the observations required special treatment and had to be justified carefully. The schedule was only be changed once submitted to MOC if there was a contingency (an instrument problem, or an operational issue that would lead to a significant loss of observing time, or for a ToO); it was not changed for "routine" tweaking of AORs.

The bottom line is that Herschel could, in normal circumstances, only guarantee to react in 7 days from an urgent request but could react in 5 days in favourable circumstances and in 4 days in a best case scenario. In theory and by taking considerable risks, faster reaction times would have been possible in a genuine emergency (e.g. a Galactic supernova), these though were never applied except in the case of operational contingencies involving instrument problems, where the need to save observing time was primordial.

Most of the steps described below applied too to normal observations too although, of course, on a much less compressed schedule.

The decision to attempt a fast turnaround time for an urgent scheduling request was not taken lightly. MOC at Darmstadt have to confirm that they could process the re-delivered observing schedule in time (even in an emergency this would not be attempted without a minimum lead time of 6 hours between delivery of a revised schedule by HSC and the start of the DTCP and, in normal circumstances, a lead time of 8 hours was required, even for a fast re-plan) and the HSC Mission Planning Group had to agree to a delivery schedule that allowed MOC enough time to carry out the full processing and check procedure for the re-planned schedule.

3.4.5. DDT requests

ToO and DDT proposals are a sub-set of Director's Discretionary Time. There were specific considerations for ToO proposals that are dealt with below. DDT proposals were submitted through HSpot following the normal guidelines for Open Time proposals; the DDT Call was always open for submissions. The most recent version of the PDFLatex proposal submission form available on the HSC Website should have been used for the submission (available here: http://herschel.esac.esa.int/Tools.shtml), although only the sections that are relevant to a DDT proposal needed to be filled in. The length of the science case needed to be commensurate with the amount of time requested (i.e. a 3 page science case was not necessary for a single 15 minute observation). An essential part of the science case was a clear explanation of why the award of Director's Discretionary Time was justified, particularly to provide a convincing case why the proposed observations should take priority in the telescope schedule over other proposals that had already been approved by HOTAC. It was recommended that a Helpdesk ticket be opened in the DDT/ToO Department confirming submission of a DDT proposal for consideration to accompany submission through HSpot, as all communication about the proposal was be made via Helpdesk, so that the history of the request and its processing was recorded and available to everyone involved in its scheduling.

As for a normal Open Time proposal, a DDT proposal could be retrieved after submission, revised and re-submitted as often as is necessary. If a proposal was revised, the proposer was advised to notify the Project Scientis and the HSC, via Helpdesk, when the final version had been submitted, so that processing of the proposal could begin.

The Project Scientist would not start to evaluate a DDT proposal until all technical reports on the contents of the proposal had been received (visibility, feasibility, AOR design, etc). When he was satisfied that he had all the information necessary for his evaluation, he would contact the HOTAC Chair and a joint decision would be made on acceptance, although the Project Scientist reserved the right to consult relevant experts, if necessary, to reach this decision. In general a decision on acceptance of DDT proposals would be taken within 3-4 weeks unless urgent scheduling considerations were involved.

For any Director's Discretionary Time proposal the main criterion for approval, if the observations were found to be feasible, was "helium into science". The proposal had to demonstrate that it would produce high enough quality science to be scheduled ahead of other, previously approved observing proposals.

3.4.6. Target of Opportunity (ToO) requests

3.4.6.1. How was a ToO alert triggered?

There were two methods. Our preferred method was "pre-approval", but ToO alerts could be triggered without pre-approval if the science case is compelling, the circumstances justified it and there was a demonstrated need to react quickly to avoid missing a major scientific opportunity.

3.4.6.2. Pre-approval by HOTAC

This was our preferred method. A normal proposal was submitted in the regular Calls for Proposals and evaluated. The proposal should have included clear trigger criteria and an expected reaction time that allowed its feasibility to be assessed. If HOTAC approved the request, the HSC was required to carry out the observations if the trigger criterion was met and scheduling constraints permitted it.

> **Note**
> Sometimes a request was simply not be feasible for operational reasons, such as having the wrong instrument active, or because of operational constraints.

This method had various advantages, not least of which is that the PI knew in advance that the observations were approved and would be made if possible and would be propietary. The observations were also already available in the HSC database, technically checked and only needed to be activated, saving valuable time on activation.

If the PI was satisfied that the trigger conditions had been met and that the target was visible, he or she sent a Helpdesk ticket in the ToO Department, supplying all the necessary information to allow the observations to be taken and a short justification of the trigger. The box stating that the ToO is pre-approved by HOTAC was ticked and the proposal name given to identify it. When the ticket was sent an SMS message went direct to the Project Scientist and to the HSC, alerting them of the triggering.

Provided that the Project Scientist was satisfied that the trigger conditions had been met, the observations would automatically be scheduled at the first possible opportunity.

3.4.6.3. Spontaneous ToO alerts

Not all ToOs could be anticipated. Sometimes something would happen that was too good an opportunity to miss; it may have been a newly discovered comet that would become bright in 6 months time, or a sudden and unexpected outburst of a known object. In this case things could be more complex and may be slower.

The PI was requested to fill out the same form in the http://herschel.esac.esa.int/esupport/ ToO Department. The alert had to be justified and the required observations either provided or, at very least in the case of extremely urgent alerts, described in enough detail that an expert at HSC could prepare them. When the ticket was sent, an SMS message went direct to the Project Scientist and to the HSC, alerting them of the triggering. No action to schedule observations was be taken at HSC until the Project Scientist had approved the request, but a fast assessment would always be made of a new proposal to get an idea of what turnaround cycle in approval and scheduling would be required, including an assessment of when the observations could potentially be scheduled; where very urgent observations were requested, the approval cycle would be accelerated as much as possible.

The Project Scientist's first reaction to a ToO request was usually to request a technical assessment of the observations and target visibility from HSC. This may have required some backwards and forwards iteration with the PI to get the observations right, so that the time impact could be assessed and the feasibility ensured. If the request was for observations of a Solar System Object (SSO) it may have been necessary for software support at the HSC to add the SSO ephemeris to HSpot. The Project Scientist sought the approval of the HOTAC Chair to add the observations to the schedule. Approval (or not) was then communicated by the Project Scientist to the PI in a reply to the Helpdesk ticket.

A ToO request could also be submitted directly via HSpot in the same way as a normal proposal, which helped in its rapid assment as it ensured that all documentation was immediately available in standard format and the proposal was in the HSC database. HSpot (since HSpot 6.2.0) generated an both and e-mail and an SMS alert of a submitted proposal, so the On Call staff at HSC and Project Scientist were aware of submission, even out of normal working hours. We recommended strongly that users send also a Helpdesk ticket to the ToO/DDT department notifying the HSC of submission when using this method.

The scientific justification for a ToO request had to be completed using the standard PDFLatex form as for normal OT2 proposals, giving a level of detail commensurate with the amount of time being requested and only filling in those sections that were relevant to the request. This form and the instructions for completing it are available here: http://herschel.esac.esa.int/Tools.shtml.

3.4.7. Processing an urgent scheduling request

Once the observations were approved by the Project Scientist, they were processed. The target day for execution was identified in consultation with MOC. Observations were normally submitted to MOC 18-20 days in advance of execution and the scheduled drafted 3 weeks or more in advance of execution. This means that any request for an observation less than 3 weeks ahead required re-processing at HSC and, usually, at MOC too.

Observations already in the database -- i.e. pre-approved ToO requests -- would be linked to the correct instrument control software and time estimator version and released to the Mission Planners, who then proceeded to de-commit the observing schedule for the day so that it could be modified. Once a de-commit of a schedule occurred it was removed completely from the system, meaning that even if the same schedule were subsequently re-instated, it had be completely re-processed at MOC, hence a schedule is only de-committed when it was certain that a re-plan is necessary and feasible and could be processed at MOC in time for uplink.

Once approved, observations not in the database had to be submitted as a new proposal first to get them into the Operational database. This became a much quicker and simpler process if they had already been submitted to HSpot by the proposer, as this made them available in the Astronomer database (the Astronomer database could be modified by PIs and their authorised co-Users, as well as by HSC staff, whereas the Operational database could only be modified by HSC staff - the Operational database was a "clean" database and only included observations previously approved for execution).

The new observation(s) were fitted into the schedule and a revised schedule generated. The revised schedule went through a list of checks and was then submitted to the instrument team(s) and the Project Scientist for final approval; either could request revision to the schedule as drafted if they identified any problems or issues with it. Only when all parties are satisfied would the schedule be transmitted to MOC. The change in schedule modified the sequence of pointing commands for the spacecraft and so had a knock-on effect for the following day too because it affected the initial state of the spacecraft reaction wheels for the next observing day, meaning that MOC had to generate a new solution for obtaining the desired pointing for both days, even though HSC had only re-generated the schedule for the first of them.

The time required at HSC to process the observations and generate the new schedule depended on the complexity of the changes to the schedule and the availability of the specialists to approve the draft schedule (out of normal working hours this was normally on a "best efforts" basis) and any modifications requested: on occasions it required a full working day and, in difficult cases, even longer to complete.

When the new schedule was approved, it was transmitted to MOC by FTP. Initial processing was automatic. Once this was complete a member of the Flight Dynamics Team would continue processing interactively. If all the checks were passed (this was not a trivial exercise, particularly with observations close to spacecraft limits), the schedule was then passed on to a member of the MOC Mission Planning Team who generated the file containing all the telecommands that would be transmitted to the spacecraft. This was usually the slowest and most complex part of the processing chain due to the serious potential consequences of any error. This processing at MOC required a full working day and had be completed in advance of the DTCP; if the DTCP was early in the day, or at the weekend, this meant that the task had to be completed on the previous working day.

Essentially, the peculiarities of the Herschel mission compared to on-line missions and the series of effectively uncompressable response times at different stages in the processing chain provided a fundamental limit to reaction times in any re-planning scenario.

Chapter 4. Space Environment

This section will deal with "space environment" aspects of the mission that affected the noise level of data and therefore the observatory sensitivity. These include:

- Background, including the telescope, instruments and the celestial background
- Radiation environment (high-energy particles)
- Source confusion (CFIRB and cirrus spatial structure, resolved or partially resolved galaxies)
- Straylight due to sources inside or outside the FoV and to instrumental self-emission

4.1. Background radiation

4.1.1. Telescope background

The Herschel telescope was located outside the cryostat and protected by the sunshade from direct radiation from the Sun. The measured telescope temperature was in the range 83-90 K and showed an annual variation due to the changing heliocentric distance of the Earth and thus L2 (see below). At this temperature, even given a low emissivity, the source contribution is almost always only a small fraction of the telescope background. For comparison, the telescope background flux was of the order of 1000 Jy, while that of Uranus was ~ 250 Jy and Neptune ~ 100 Jy. Therefore, a precise characterisation of its behaviour is of critical importance.

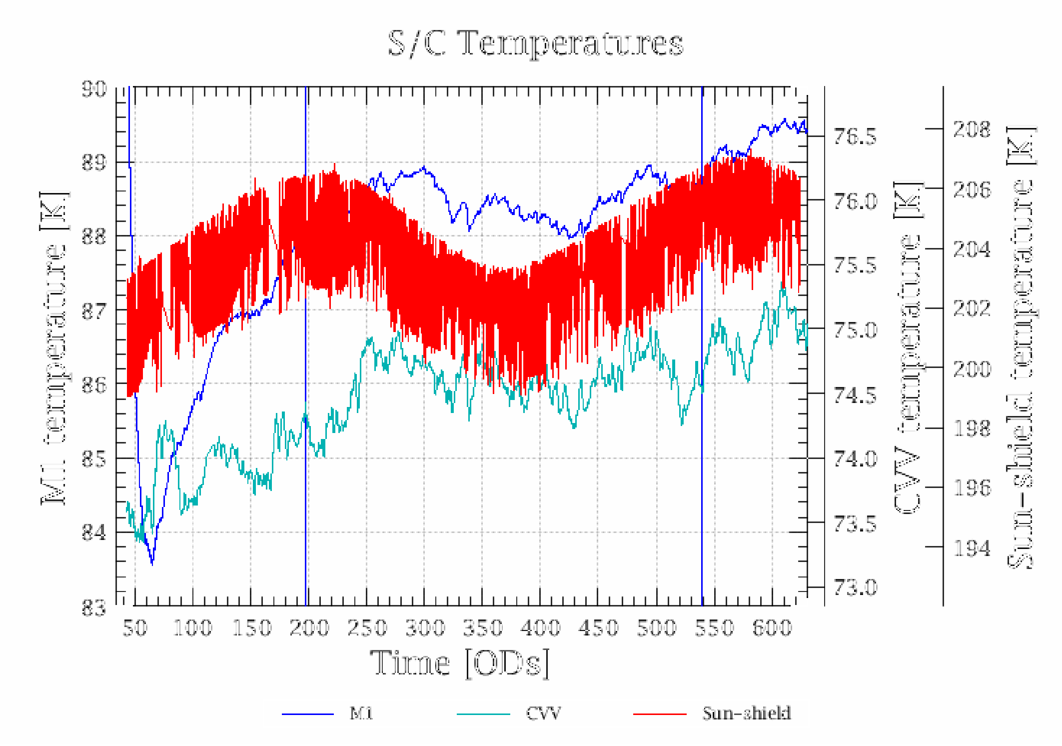

Figure 4.1. Temperatures of the primary mirror (M1), cryostat vacumm vessel (CVV) and sun-shield measured from OD40 to OD660. The monotonic increase of temperature from up to OD300 was well correlated to the seasonal temperature variation model.

The telescope background depends primarily on:

- The average temperature: Figure 4.1 shows the temperature of Herschel's telescope primary mirror measured across 620 ODs (from 22nd July, 2009 to 4th March, 2011; the CVV and sunshield temperatures are also displayed). The data gathered indicate that the temperature was some 5K higher than predicted pre-launch (this was compensated by the lower telescope emissivity, see below). On the other hand, the overall trend, characterised by a monotonic temperature increase with a maximum amplitude of some 6K, is well correlated with the 'seasonal' temperature evolution mode.
- The effective emissivity: beyond 100µm, this has a stronger influence on the telescope background level than the temperature. It has been observed that a 1% reduction in emissivity gave a greater improvement than a 5 K reduction in temperature. The predicted (modelled) telescope emissivity was < 0.4%/0.8% (for clean/dusty samples; Fischer et al., 2004 [RD11]). Preliminary results from PACS photometer observations suggest that the actual telescope emissivity was quite significantly lower (about half the predicted figure; Okumura, priv. comm.)
- The straylight (see Section 4.4).

Small spatial and temporal temperature gradients are important to the background stability. The requirements on the primary mirror (M1) are:

- Maximum temperature difference along the S/C Z axis < 10 K (predicted < 0.5 K)
- Maximum temperature difference along the S/C Y axis < 1 K (predicted ~ 0.0 K)
- Along the S/C Z axis: dT/dt < 13.0 mK/min
- Along the S/C Y axis: dT/dt < 1.3 mK/min

4.1.2. Instruments

See "Self-emission" under Section 4.4

4.1.3. Celestial background

Thermal emission from interstellar dust (known as interstellar cirrus) dominates the FIR Sky Background (FIRSB) at lower Galactic latitudes, while the Cosmic Far-Infrared Background (CFIRB) is more significant towards higher Galactic latitudes, also dominating the confusion noise in the PACS and SPIRE photometric bands. Intrinsically diffuse and unresolved components of the FIRSB are (in descending order of their relative contribution, see [RD4] and [RD5]):

1. Diffuse galactic light (interstellar cirrus): quasi-thermal emission of dust in low density gas clouds in the Milky Way. This is the dominant component for wavelengths λ>70µm.
2. Zodiacal light and emission from the asteroid belt: this is the dominant component of the sky background at MIR wavelengths.
3. Cosmic far-infrared background (CFIRB): accumulated and unresolved light of distant galaxies.
4. The cosmic microwave background (CMB): the CMB also has an important contribution in strength, but the fluctuation amplitudes are small, and well below the detection limits of PACS and SPIRE.
5. Intergalactic diffuse emission
6. Integrated starlight: the integrated contribution from faint stars in the Milky Way is an important component for near- to mid-infrared wavelengths, but has a negligible contribution for longer wavelengths, e.g. those of PACS and SPIRE, as even the coolest stars are far from the

black body peak, even at the shortest PACS wavelengths.

A detailed description of the different components of the FIR background is given in the Herschel Confusion Noise Estimator (HCNE) tool Science Implementation Document ([RD4] and [RD5]). The HCNE tool can be accessed as a standalone service to provide background estimates (see the HSC website for more information), or through the HSpot proposal preparation tool.

While the zodiacal light emission is a major contributor to the sky brightness in the MIR range, it is less important for the FIR and sub-mm wavelengths. Moreover, this emission is quite smooth, lacking fluctuations on arcmin scales (the angular resolution of the ISOPHOT instrument on board ISO). Smaller scale fluctuations, in principle, are likely to exist, but the presence of such structures have not been yet confirmed by the recent observations of the Spitzer Space Telescope.

Confusion noise due to the integrated FIR-sub-mm emission from faint asteroids, individually below the detection limit, has been investigated by Kiss et al. (2006) (see [RD5] and references therein). It has been found that the distribution of asteroids concentrates towards the local anti-solar direction, with a corresponding peak of the confusion noise in the anti-solar point, and an extended cloud is present around the maximum. Seasonal variations are also detected. The confusion noise induced by the cloud of asteroids would only be not negligible in the area around the anti-solar direction, but this area of the sky is closed to Herschel anyway due to the satellite's Sun constraint (see Section 2.3), so the asteroid cloud component is not considered in the HCNE.

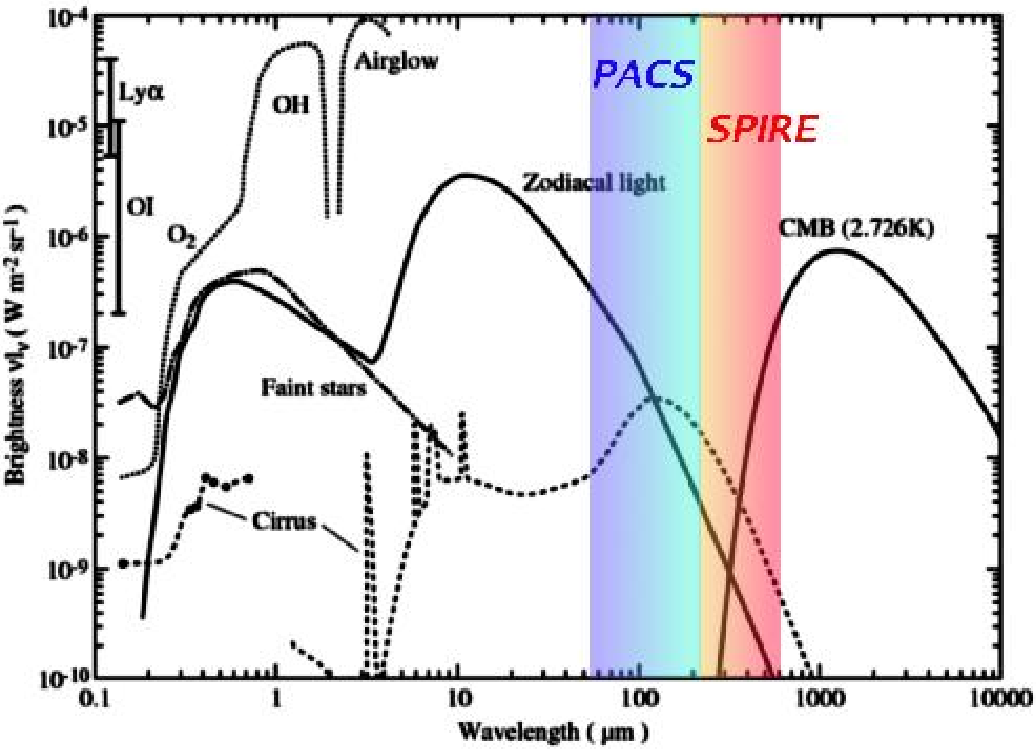

Figure 4.2. The brightness of the night sky, excluding the contribution of the extragalactic background (from [RD5], adapted from Leinert et al. 1998, A&A, 127, 1). The spectral ranges covered by the PACS and SPIRE instruments of the Herschel Space Observatory are indicated. Atmospheric contributors, affecting ground-based observation in the optical and NIR, have been also displayed. In the absence of confusion, the far infrared sky is darkest around 350 microns where zodiacal light, CMB and cirrus are all close to their minimum values.

The interstellar medium shows a strong concentration of emission around the Galactic plane; this feature is conspicuous at many wavelengths. However, the cirrus emission is not limited to low Galactic latitudes. It consists of thermal emission of dust in low-density, cool interstellar HI clouds (typically with T≈20K and n≤10^2cm^{-3}), showing a smooth, modified blackbody SED. It is a strong source of emission, and dominates the sky for wavelengths λ>70µm, even at high Galactic latitudes. The cirrus emission is highly structured, and shows a typical filamentary structure.

The main characteristic of the cirrus emission is its spatial structure at a specific wavelength. This is usually described by the spectral index, α,of the power spectrum of the image, averaged over annuli ([RD7]). With this parameter the power spectrum is $P=P_o(f / f_o)\alpha$, where P is the power at the spatial frequency f and P_o is the power at the spatial frequency f_o. Due to this parameterisation the structure of cirrus is equivalent to that of a fractal.

According to [RD5], cirrus confusion noise can be generally described by the following equation:

$$\sigma_{cirrus} = c_1 \times (\lambda/D)^{1-\alpha/2} \times B^{\eta}$$

Here σ_{cirrus} is the confusion noise due to the cirrus component, B is the surface brightness of the field α is the spectral index of the logarithmic power spectrum, averaged in annuli (see [RD5] and references therein), λ is the wavelength of the observation and D is the effective diameter of the telescope's primary mirror. The parameters c_1 and η have to be determined from measurements. This is used within the HCNE to compute the noise due to cirrus emission. Details of the computations are given in [RD5].

In many practical cases, Galactic cirrus confusion noise has been found to be easily parameterised as follows (see for instance [RD6] and references therein):

$$\sigma_{cirrus} \sim 0.3(\lambda_{100})^2(D_m)^{-2.5}\langle B_\lambda\rangle^{1.5}$$

where σ_{cirrus} is given in mJy, λ_{100} is the wavelength ratio λ/(100 µm), D_m is the telescope diameter in metres and $\langle B_\lambda\rangle$ is the sky brightness in MJy/sr. If we consider fiducial values $\langle B_{70}\rangle$ = 0.12 MJy/sr and $\langle B_{160}\rangle$ = 1.5 MJy/sr (corresponding to $N_{HI} = 10^{20}$ cm^{-2}) and D_m = 3.5, we get that σ_{cirrus}(70 µm) = 0.22 µJy and σ_{cirrus}(160 µm) = 0.08 mJy.

4.2. Radiation environment

The L2 environment (and orbits around it) is relatively benign compared to those in geostationary (GO), or low Earth (LEO) orbits. In particular, a series of common threats for satellites in GO or LEO, including the neutral thermosphere, space debris, geomagnetically trapped particles and large temperature gradients, are not a concern for L2 orbits. Environmental aspects to be considered at L2 include:

- Solar wind plasma. Essentially a neutral or cold plasma: 95% protons, 5% He++ and equivalent electrons; 1-10 particles/cm^3. The main risk associated is a low surface charging potential. This plasma may be relatively benign at L2 compared to that found at GO and LEO.
- Ionising radiation: solar flares (energetic electrons, protons and alpha particles), Galactic cosmic rays and Jovian electrons.
- Magnetic fields: Earth's magnetotail extends up to 1000 Earth's radii, so it must be considered (2-10 nT) along with interplanetary magnetic field (~ 5 nT). The effects on the spacecraft and PLM include possible orbit disturbance and electrostatic discharge (ESD).

Therefore, the main radiation components at L2 consist of: Galactic cosmic rays, solar particle events and solar and Jovian electrons.

In the early stages of the mission, the dominant radiation source was Galactic Cosmic Rays, which were at the highest level ever recorded at launch, due to the record depth of solar minimum, supplemented by Jovian electrons, characterised by a energetic population and a 13-month synodic year modulation.

The Herschel spacecraft was equipped with a Standard Radiation Environment Monitor (SREM) placed in the -Z SVM panel; the SREM is a particle detector developed for satellite applications that has been added to Herschel and Planck as a passenger. It measures high-energy electrons (from 0.5 MeV to infinity) and protons (from 20 MeV to infinity) of the space environment with an angular resolution of some 20 degrees, providing particle species and spectral information. The SREM data were received on-ground and processed by the Space Weather Group at ESTEC, providing valuable information on the radiation environment at L2. A sample plot showing the calibrated count rates in three counters (TC1 - protons with E > 20 MeV; TC2 - protons with E > 39 MeV; TC3 - electrons with E > 0.5 MeV) is displayed in Figure 4.3.

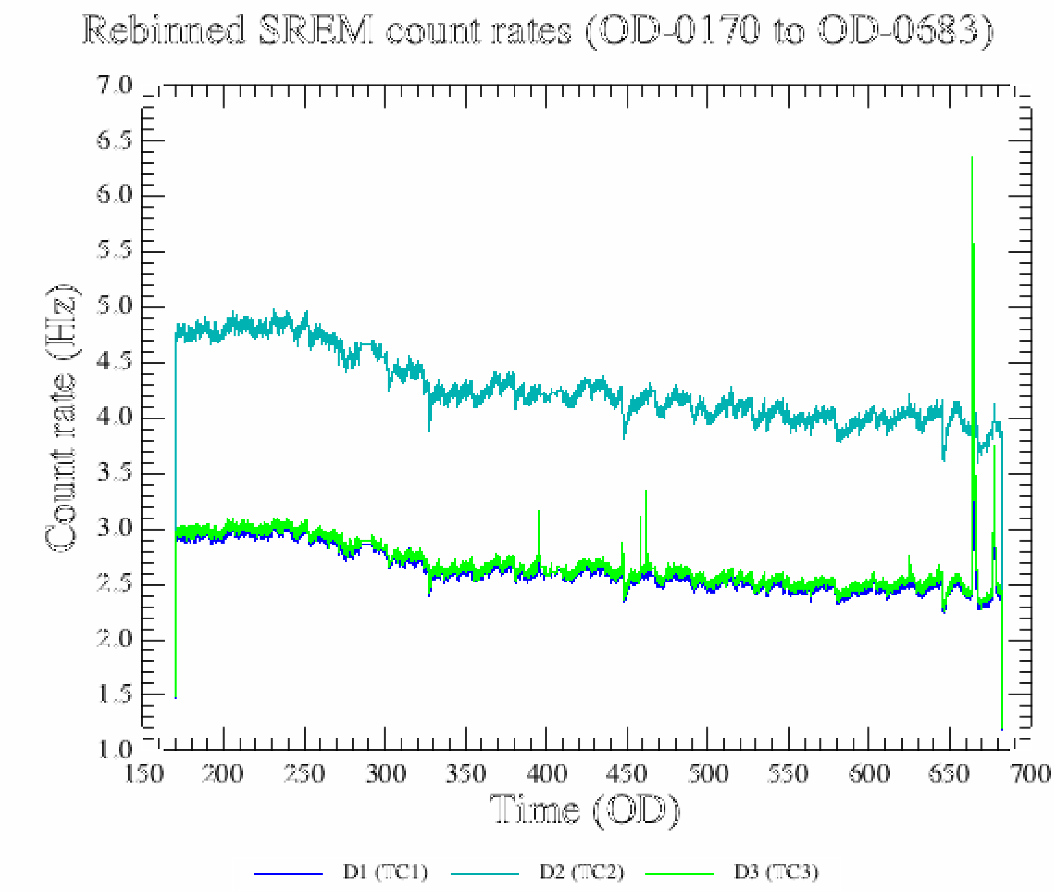

Figure 4.3. SREM calibrated count rates in three counters (TC1, TC2 and TC3), re-binned in intervals of five minutes. from the 30th of October 2009 (OD 170) to 27th March of 2011 (OD 683). The slight decline of the count rates can be explained by an increased solar activity and the subsequent increase of shielding to Galactic cosmic rays. Several events are visible, the most conspicuous, a small proton flare, detected in OD-663 (7-8 March 2011).

Weekly, calibrated plots of the Herschel SREM data and special plots of any observed proton events are available to users, provided by the SREM PI, Petteri Nieminen, using a more sophisticated analysis of the data. These are available at http://proteus.space.noa.gr/~srem/herschel/. All Herschel SREM data are subject to continuous re-processing as the algorithms used are refined and the calibration improved.

4.2.1. Solar activity and its influence on Herschel

4.2.1.1. The Solar Cycle

Solar activity follows an approximately 11-year cycle. The last minimum occurred in September 2009, with a record low sunspot number of 8.4 (http://solarscience.msfc.nasa.gov/predict.shtml) and therefore the Herschel launch in 2009 happened during an exceptionally low activity state. Contrary to initial pre-launch predictions, the current solar cycle has been well below average in intensity, with a predicted maximum sunspot number of 66 in autumn 2013, six months after End of Helium. The progress of Solar Cycle 24 up to summer 2013 is displayed in Figure 4.4, along with the predicted evolution. Solar Cycle 24 has had the lowest maximum since Solar Cycle 14, which peaked at 64 in 1906.

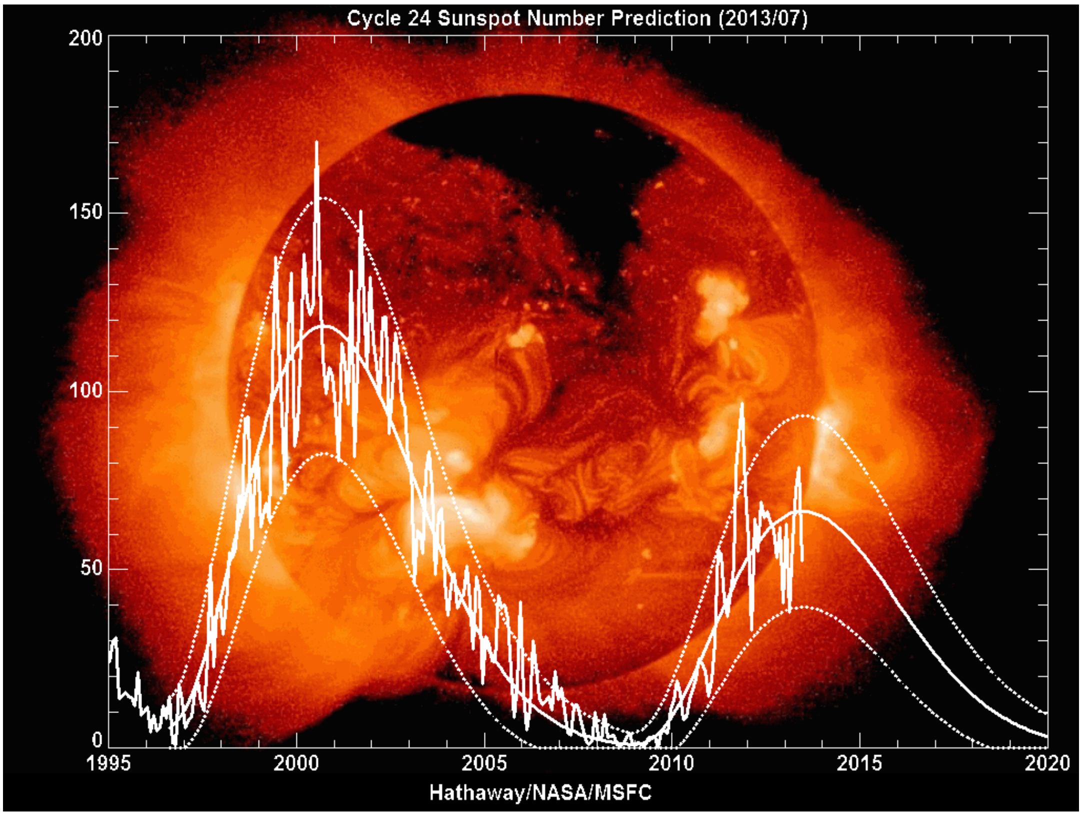

Figure 4.4. The smoothed sunspot number since 1995, covering Solar Cycle 23 and 24, along with the end-of-helium version of the predicted curve, which had undergone radical revision since the pre-launch previsions of a possibly extremely high maximum for Solar Cycle 24. In fact, the prediction for maximum has been revised downwards every year since 2009 and the current maximum is on course to be the lowest since 1906 (Solar Cycle 14) with its peak of 64.

The chief impact of high solar activity is Solar Proton Events. These are produced by major, long-duration x-ray flares close to the centre of the solar disk, or by flares that may be close to the solar limb, or even on the farside of the disk, but that are connected to the Earth by the spiralling interplanetary magnetic field. A typical event, observed on January 23rd 2012, is shown in Figure 4.5, showing the relative timing of the x-ray flare, the arrival of the high-energy protons and the later shock wave impact when the Coronal Mass Ejection arrived, bringing a further burst of, principally, lower energy protons. In contrast to a normal x-ray flare, which has a duration of a few minutes from start to finish, a long-duration flare may be active for 12 hours or more and have an integrated energy two orders of magnitude greater than a normal flare. The proton flare may be active for several days, but the energy spectrum gets softer rapidly as the high-energy protons fall off much more quickly than the low-energy protons.

X-ray flares and Solar Proton Events are correlated with solar activity, but not in a simple fashion. A large sunspot number does not necessarily correlate with high levels of proton activity. Major flares occur when there are sunspot groups with highly complex magnetic fields that contain large amounts of magnetic energy. Solar Proton Events are rare close to solar minimum (none were observed in 2007-2009 and only a single, very small event, in 2010), but may occur at any time for several years around solar maximum; the maximum of solar proton activity may not coincide with the peak of the sunspot cycle.

4.2.1.2. Solar proton events and the cosmic ray flux

With solar maximum expected late in 2013, solar particle events were expected to be problematic only towards the end of the mission. Solar activity increased steadily after launch, but solar proton activity was modest and much lower than in Cycle 22 or 23.

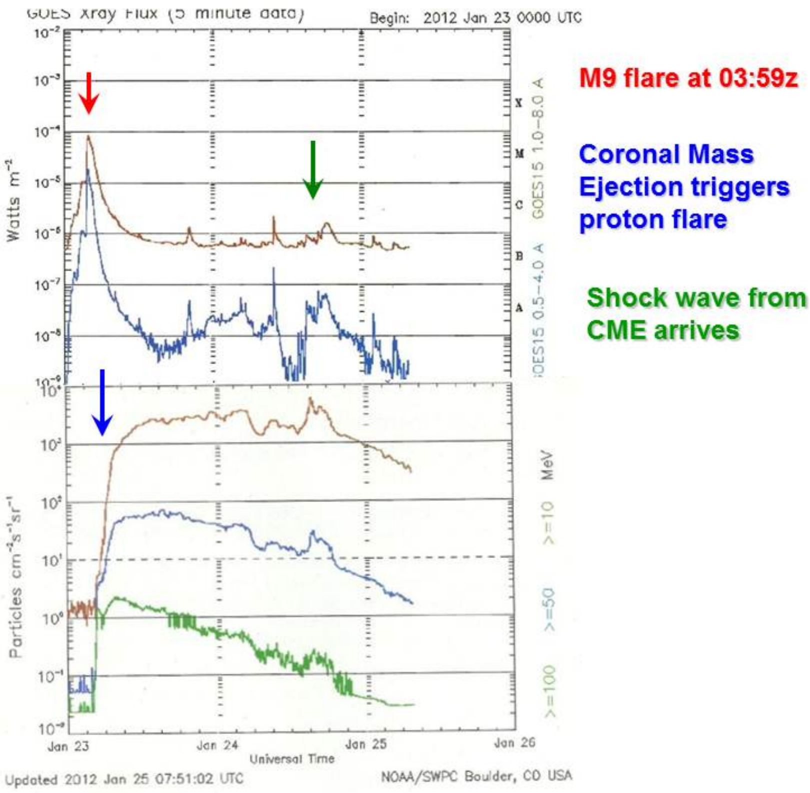

Figure 4.5. The sequence of x-ray (top) and high-energy proton (bottom) activity. From top to bottom we have the soft x-rays (1-8 Angstroms), hard x-rays (0.5-4 Angstroms), low energy protoms (>10MeV), medium energy protons (>50MeV) and high energy protons (>100MeV). Evidence suggests that Herschel was only sensitive to very high energy protons.

> **Note**
> The amplitude of solar proton storms is measured in "proton flux units" (pfu) at >10 MeV, where 1 pfu is 1 proton cm^{-2} s^{-1} sr^{-1}. The threshold for a storm is 10 pfu. In the literature the storm intensity is often given as the integer part of the logarithm of the peak flux in pfu (i.e. S1 intensity is 10-99 pfu, S2 is 100-999 pfu, etc. Solar Cycle 23 featured six major storms of S4 intensity, although none has yet occurred in the current Solar Cycle 24.

Two proton events of around 6000 proton flux units observed in January and March 2012 were the strongest ones observed to End of Helium (as of December 2013 they are still, by some distance, the strongest events of this current maximum). In contrast. the previous solar maximum gave three events of greater than 20000 proton flux units in 2000, 2001 and 2003. Solar Cycle 22 produced events over 40000 proton flux units in 1989 and 1991.

Notably, there is only a limited impact of this increased solar activity in the performance of the instruments or the quality of the science products at the levels of activity observed so far. Unlike many other astronomical satellites, Herschel has proved to be highly resistant to the effects of solar storms. Higher levels of proton flux increase the observed glitch rate and background offsets due to charge transfer from proton hits changing the effective detector biasing, most notably in PACS spectrometer observations. While chopped modes are relatively robust against such effects, the unchopped PACS spectroscopy mode was particularly vulnerable to possible degradation during solar proton storms.

> **Note**
>
> After a proton storm, any data that were potentially affected were carefully assessed in the Quality Control process, including an inspection by eye by one of the HSC instrument experts. No observation has been failed in this Quality Control process during the mission.

A potential issue that was, in the end not seen in data, was the failure of observations due to over-biasing from charge transfer at elevated glitch rates. As each PACS spectrometer observation was performed using a matching capacitor, tuned to the expected flux, there was a danger that a high glitch rate could bias the detector to such a point that it would saturate on the background, making the data useless. The danger level could only be determined empirically. The only statement that we can make is that no observations were failed by over-biasing up to a level of Solar Proton Storms of 6000 pfu but, it is perfectly possible that at 6500 pfu, observations could have failed.

In Figure 4.7 the more than four years of data between launch and the final passivation of Herschel are presented on a single plot for easier direct comparison and to make the variation in the cosmic ray flux -- the base level for the plots -- clear. The most obvious feature is the progressive decline in the cosmic ray flux starting early in 2010 as the heliosphere expanded outwards with increased solar activity and served to attenuate the cosmic ray flux arriving at L2. This cosmic ray flux reached record levels during the recent, unusually deep and prolonged solar minimum, when the measured flux was over 20% higher than at any previous time since observations began in 1958.

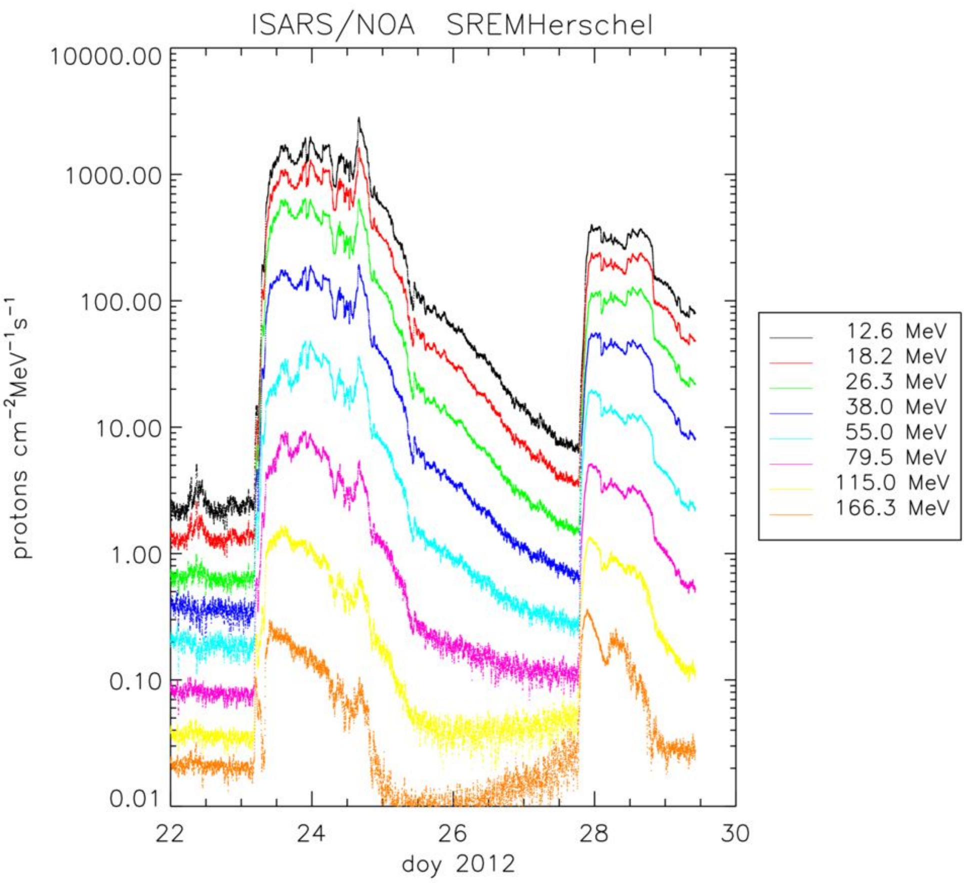

Figure 4.6. SREM calibrated count rates for the 2012 January 23 and January 28 solar proton storms. The arrival of the shock wave from the Coronal Mass Ejection can be seen as a sharp peak in the proton flux approximately 36 hours after the initiation of the January 23 event; this peak is most clearly seen in the data at the lowest energies. For some of the Solar Proton Events observed during the Herschel mission, due to the rapid decay of the event at higher energies, there may be little or no signal at energies above 50MeV when the shock wave finally arrives.

The solar proton events observed in Solar Cycle 24, as seen increasingly in the SREM data through 2011 and 2012, are soft, with the amplitude dropping rapidly to higher energies. However, events differ widely in energy spectrum and so may not be comparable. The reference energy for solar proton events is the 10 MeV flux (see: http://umbra.gsfc.nasa.giv/SEP/ for a listing of all solar proton events since 1976, with some useful background information, although this listing does not include the very largest event ever registered, that of November 1972, which occurred just 3 days after Apollo XVII had landed and which, in all probability, would have been fatal for the astronauts had it occurred during the flight) but, for example, although the events of 2012 January 23 (Figure 4.6) and 2012 March 7 were of similar amplitude at the reference energy of 10 MeV, the latter was approximately a factor of 6 larger at 166 MeV. For geoeffective Coronal Mass Ejections, in which the Coronal Mass Ejection is aimed at the Earth (this is usually the case when the sunspot is close to the centre of the solar disk), the peak 10 MeV flux is usually measured when the Coronal Mass Ejection shock wave hits, typically around 24-36 hours after the start of the event; this shock wave may have only a very low amplitude at energies above 50 MeV and occur when the flux of higher energy protons is already declining rapidly, so it will often be weak or even absent at high energies, for which the highest proton flux is measured during the initial large proton burst.

We find no evidence that there is a correlation between solar proton storms and instrumental Single Event Upsets (SEUs) -- bit flips in the memory -- on board the spacecraft. All evidence suggests that instrument SEUs are linked to impacts from high energy cosmic rays, well above the energies registered typically in the solar proton storms of Solar Cycle 24, although very hard solar events have been registered in previous solar maxima that could potentially have had an effect; the hardest proton event on record was registered in 1989, during the maximum of Solar Cycle 22. However, glitch rates in the PACS spectrometer do increase during Solar Proton Events, with the March 2012 storm giving an approximately 70% increase in glitching; as glitch rates increase, the main uncertainty in dtermining their rate is the difficult in determining what is actually a glitch in the data as confusion between glitches can become important. At the same time, the impact of high energy particles in the detectors leaves charge, causing an effective increase in the bias level and sensitivity changes that must be calibrated out in processing.

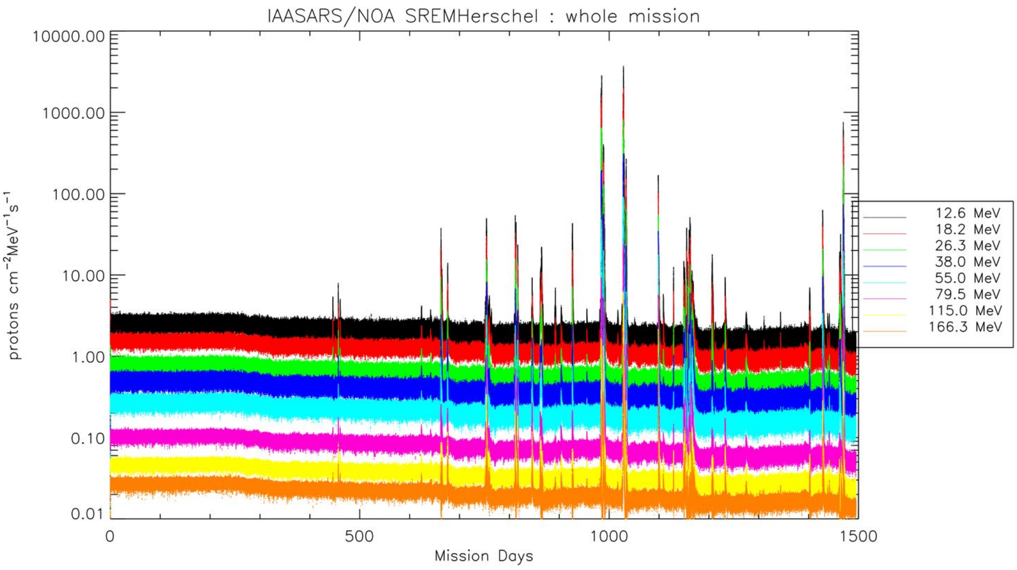

Figure 4.7. SREM calibrated count rates from the archive of calibrated SREM data http://proteus.space.noa.gr/~srem/herschel/ for the entire mission. There was a higher level of background cosmic ray flux at the start of the mission, corresponding to solar minimum. From early 2010 the cosmic ray flux dropped continuously to the end of mission. There are three very small gaps in the plotted data due to the fact that on three ODs early in the mission, on-board anomalies led to corrupted SREM data. Plot prepared by Ingmar Sandberg on behalf of the SREM Team.

4.2.1.3. SEUs

Solar Proton Events and cosmic rays are known to be associated with bit flips - Single Event Upsets, or SEUs - in computer memory. Charged particle impact in memory causes the value of the impacted bit to change, corrupting the memory. While, in most cases, this is inoffensive, hits in critical memory areas could cause an instrument to enter an unstable condition, producing data that might be corrupted, or even garbage, or even send the instrument directly into a safe mode, as happened on several occasions during the mission. The failure of the primary HIFI power chain on OD-86 was finally traced to a sequence of events, including a power surge, triggered by an SEU in a critical area of memory.

Studies have shown that there is no link between epochs of Solar Proton Events and increases in the rate of SEUs in Herschel. Between HIFI being switched back on in early 2010 and EoH, the median interval between HIFI SEUs was 10.5 days while, for SPIRE, it was 42 days. However, between 2011 July 10 and 2012 March 7, when the Sun was most active in terms of SPEs, there were just 3 HIFI and 2 SPIRE SEUs in an interval of 241 days. In contrast, in 2010, when there was only a single SPE that barely crossed the threshold of 10pfu, both SPIRE and HIFI registered the largest number of SEUs in any year of the mission. The lack of an obvious increase of SEUs at times of increased SPE activity leads us to believe that the most likely source of SEUs is high-energy particles, with energies probably >500MeV, of cosmic origin. This conclusion is reinforced by the obvious trend to lower rates of HIFI SEUs through the mission, albeit in noisy data caused by small number statistics, closely following the trend in the cosmic ray flux measured by the SREM.

Herschel's on-board mass memory also suffered bit flips that were detected and corrected. Discounting a few events that were multiply detected and thus cause considerable noise in the statistics (all occasions where more than ten bit flip corrections have been recorded in a single day are due to multiple detections of single events -- these can be identified by their regular spacing in the data record due to the fact that the on-board SEU check and correction routine ran at regular intervals), the mean frequency of Mass Memory SEUs was 1.0 per day through the mission although, during the first six months after launch, the rate was approximately 30% higher (see Figure 4.8). From the start of 2010 to the end of the mission, the rate of Mass memory bit flips remained constant within the errors (see Figure 4.9) with no trend in the rate; however, during this time the cosmic ray background as measured by the SREM dropped by a factor of approximately 2 (Figure 4.7), a reduction that is not reflected in the Mass memory bit flip rate.

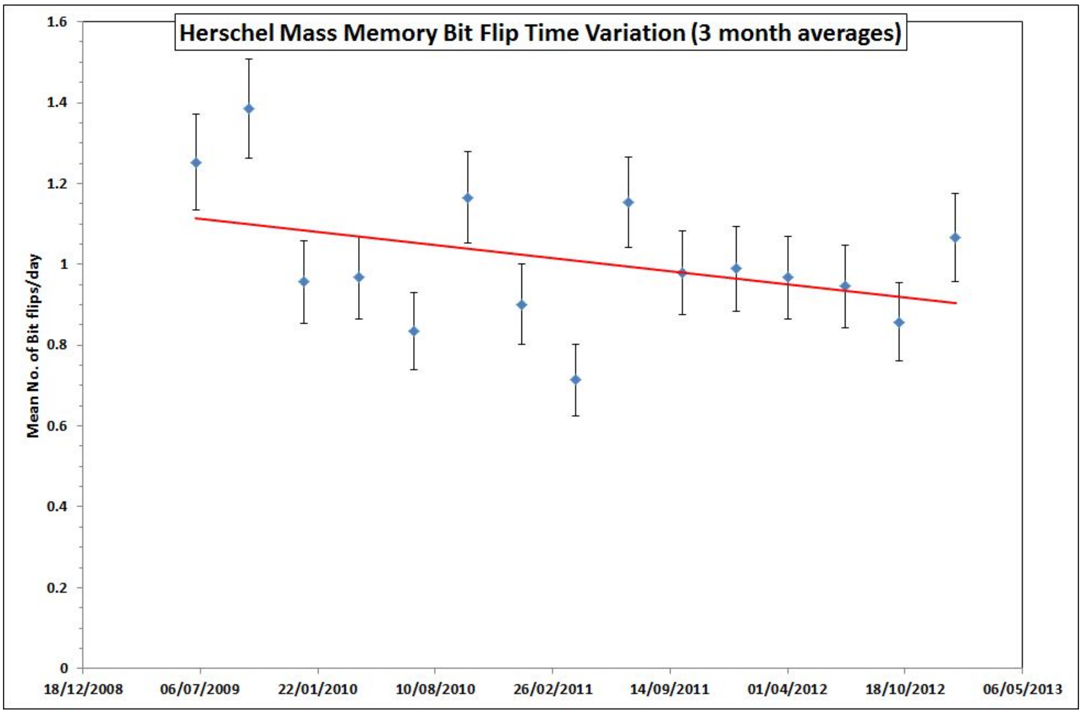

Figure 4.8. The daily rate of bit flips in Herschel mass memory through the mission, averaged over 3 month intervals, along with the best least squares fit to the data. Plot prepared at HSC from information supplied by MOC.

It is not obvious why the rate of Mass Memory bit flips does not trend with the background rate of cosmic rays. A CERN study by Stozhkov et al. "Cosmic ray measurements in the atmosphere" shows that the cosmic ray flux from 40MeV -- 2.4GeV, at 31km in the Earth's stratosphere, between 1957 and 2001, is strongly dependent on the solar cycle, with a range in cosmic ray flux of approximately 50% at the highest energies and a factor of 2 or more at the lowest energies in this range (which are typical of the energies detected by the Herschel SREM). The correlation coefficient for this dependence is --0.82+/--0.03.

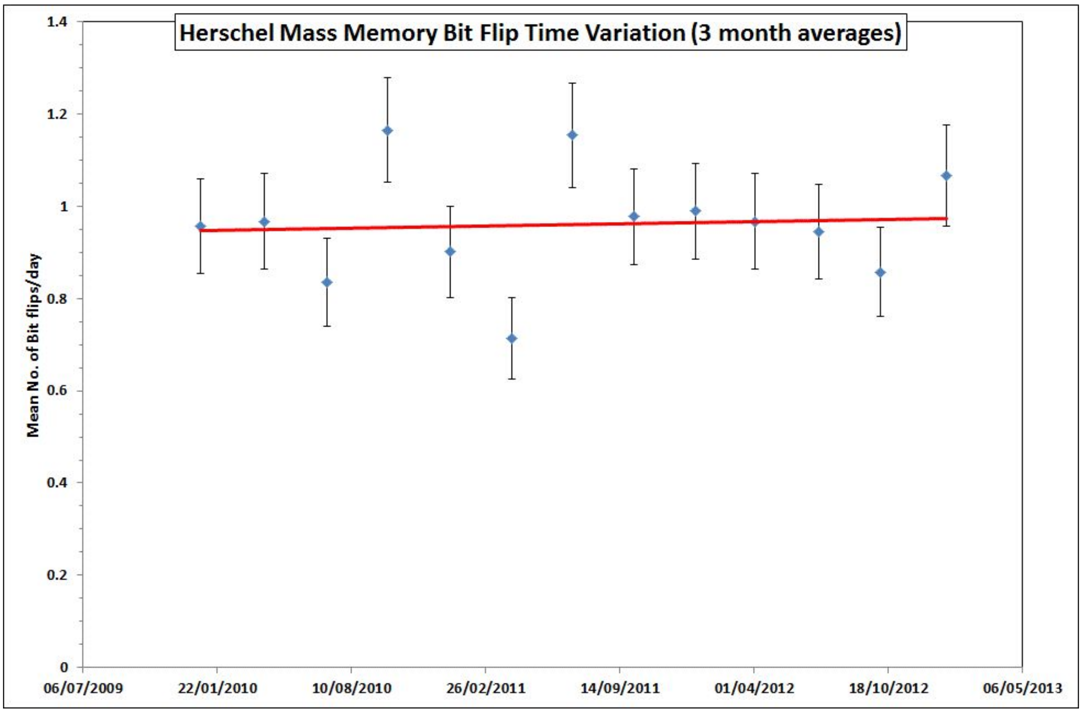

Figure 4.9. As in Figure 4.8, supressing the data from the first six months of the mission when the rate of Mass Memory bit flips was highest (this conforms to a statistical rule of thumb that "if a trend in a graph disappears when you cover about 10% of the data with your thumb, then the trend is, most likely, not real"; the rate of bit flips is seen to be constant, within the errors -- the best least squares fit, in red, shows no significant trend in rate -- and the dispersion is consistent with a Poissonian distribution of errors in a constant signal. Plot prepared at HSC from information supplied by MOC.

4.3. Source confusion

Source confusion is an additional noise factor closely related to the astronomical background, described in Section 4.1. The sensitivity limit due to confusion is determined by the telescope aperture, observation wavelength and the position on the sky. The sensitivity cannot be improved by increasing the integration time after reaching the confusion limit. The most important contributions to source confusion are:

- Structure of the CFIRB, as well as resolved and partially resolved extragalactic sources dominate at high galactic latitudes.
- Small-scale structure in cirrus clouds may dominate at intermediate Galactic latitudes. The contribution depends heavily on the level of cirrus emission at the position on the sky.

The confusion noise is usually defined as the (stochastic) fluctuations of the background sky brightness below which sources cannot be detected individually. In addition to the diffuse Galactic foreground cirrus component, these fluctuations are caused by intrinsically discrete extragalactic sources in the beam. A classic example of source confusion is seen in the SPIRE image of the Lockman Hole Figure 4.10, in which, over the image, the background is defined only by the fainter, unresolved sources. A confused field in which, additionally, there is confusion due to cirrus emission is shown in Figure 4.11. Due to the limited telescope diameter compared to the wavelength, these fluctuations play an important, if not dominant, role in the total noise budget in extragalactic surveys carried out in the MIR, FIR and sub-mm range. Moreover, the noise due to extragalactic sources depends strongly on the shape of the source counts at a given wavelength.

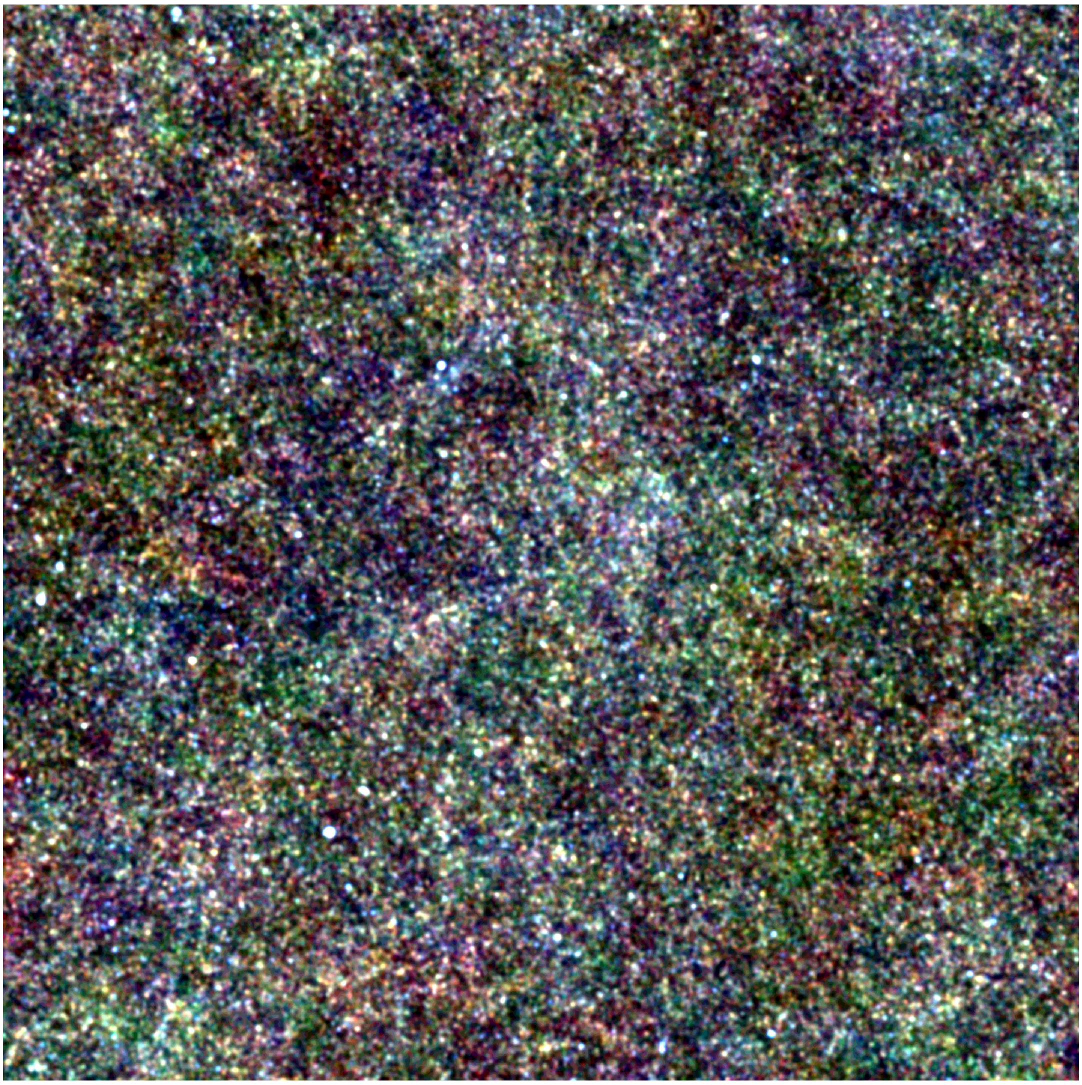

Figure 4.10. An example of source confusion. This SPIRE image of the Lockman Hole (the red layer is 500 microns, green is 350 microns, blue is 250 microns) the surface density of galaxies is so great that the entire image is filled. The background is not dark sky, but instead it is composed of fainter and more distant galaxies on which the brighter galaxies are superimposed. See Figure 3.2 for an explanation of the significance of the colours of the different sources.

There are two different criteria to derive the confusion noise, and thus the detectability of a point-like or compact source:

- First, the target source flux should be well above the average background fluctuation amplitude. This is the basis of the "photometric criterion", derived from the fluctuations of the signal due to sources below the detection threshold S_{lim} in the beam.
- On the other hand, the observed source should be far enough from its neighbours to be properly separated; this is the basis of the "source density criterion", which is derived from a completeness criterion and evaluates the density of the sources above the detection threshold S_{lim}, such that only a small fraction of the sources are missed because they cannot be separated from the nearest neighbour.

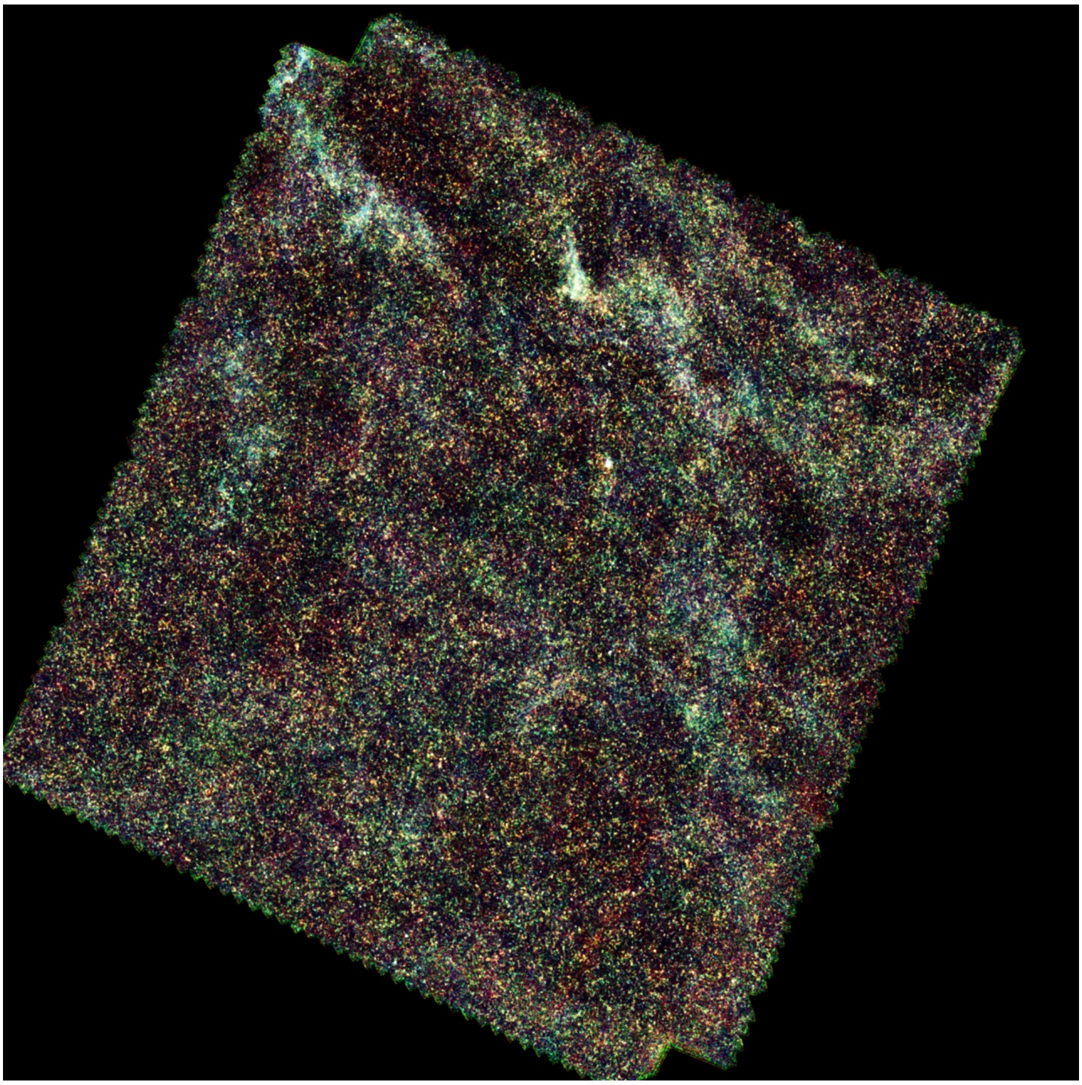

Figure 4.11. An example of source confusion compounded by cirrus emission. In this section of the ATLAS survey a region of cirrus emission can be seen near the top of the image, adding to to the confusion caused by the high density of sources.

Generally, we should compare the confusion noise derived from both criteria, in order not to underestimate it artificially. The confusion noise, σ_c, and confusion limit, S_{lim} are defined as follows:

There are two different criteria to derive the confusion noise, and thus the detectability of a point-like or compact source:

- First, the target source flux should be well above the average background fluctuation amplitude. This is the basis of the "photometric criterion", derived from the fluctuations of the signal due to sources below the detection threshold S_{lim} in the beam.
- On the other hand, the observed source should be far enough from its neighbours to be properly separated; this is the basis of the "source density criterion", which is derived from a completeness criterion and evaluates the density of the sources above the detection threshold S_{lim}, such that only a small fraction of the sources are missed because they cannot be separated from the nearest neighbour.

Generally, we should compare the confusion noise derived from both criteria, in order not to underestimate it artificially. The confusion noise, σ_c, and confusion limit, S_{lim} are defined as follows:

$\sigma_c{}^2 = \int f^2(\theta,\phi)d\theta d\phi \int_0^{S\_lim} S^2(dN/dS)dS$

where $f(\theta,\phi)$ is the instrumental 2D beam profile, that can be approximated by a Gaussian profile with the same FWHM as the expected PSF, or by an Airy function, S is the source flux density (in Jy) and dN/dS is the differential source number counts (in $Jy^{-1}sr^{-1}$).

Then, the total noise is computed by adding in quadrature the different noise contributions, in this case the photon (and instrumental) noise and the confusion noise, i.e. $\sigma_{total} = (\sigma_p{}^2 + \sigma_c{}^2)^{½}$

The photometric criterion is defined by choosing the S/N ratio q_{phot} between the faintest source (of flux S_{lim} and the noise σ_c due to fluctuations from beam to beam caused by sources fainter than S_{lim}, as given by the implicit equation:

$q_{phot} = S_{lim} / \sigma_c(S_{lim})$

q is usually chosen between 3 and 5, depending on the specific objectives.

The source density criterion is defined by setting the minimum degree of completeness of the detection of sources above the limiting flux S_{lim}, which is driven by the fraction of sources lost in the detection process due to a nearest neighbour source with flux above S_{lim} too close to be separated given an instrumental beam size. For a given Poissonian source density N(>S), the probability P of finding a nearest neighbour with $S \geq S_{lim}$ at a distance closer than the minimum angular separation θ_{min} is given by:

$P(< \theta_{min}) = 1 - \exp(-\pi N\theta^2_{min})$

An acceptable probability limit is $P_{max} = 0.1$. The minimum distance is usually parameterised using the FWHM of the beam profile $\theta_{min} = k\theta_{FWHM}$, and $0.8 \leq k \leq 1$. Fixing the probability we obtain the corresponding "source density criterion" limiting density of sources:

$N_{SDC} = -\ln(1-P(<\theta_{min})) / \pi N k^2\theta^2_{FWHM}$

The instrumental beam area, is given by $\Omega \sim 1.14\theta^2_{FWHM}$. Therefore, for P = 0.1 and k = 0.8, the density is 1/16.7 sources/beam. The limiting source flux, S_{SDC} is thus determined by using existing number counts results and a suitable model for infrared galaxy evolution extrapolating the data to the appropriate wavelengths and (faint) flux levels. The confusion noise, σ_{SDC} is computed using the same relation as for the photometric criterion, as the S/N ratio $q_{SDC} = S_{SDC} / \sigma_{SDC}$.

The Herschel confusion noise levels due to extragalactic sources in the different instruments/bands are being computed from deep maps on 'blank' fields. The preliminary results ([RD11] and [RD12]) are shown in Table 4.1 along with the values predicted by by Lagache et al. 2003 ([RD8]) using number counts derived from a phenomenological model based on template spectra of starburst and normal galaxies, and on the local infrared luminosity function. This model has been found to be in very good overall agreement with ISOCAM at 15 µm, IRAS at 60 and 170 µm and SCUBA at 850 µm (see references within [RD8]). Confusion level predictions for Herschel/PACS have been also computed by Dole et al. 2004 ([RD9]) based on recent Spitzer/MIPS number counts from Papovich et al. 2004 ([RD10]) shown in Figure 4.12. They obtain S_{SDC}(70 µm) = 0.16 mJy and S_{SDC}(160 µm) = 10.0 mJy. Please refer to the specific instruments' Observers' Manual for an up-to-date information regarding confusion noise.

Table 4.1. PACS and SPIRE measured confusion noise, compared to predictions computed according to photometric and source density criteria. From [RD9]. I.

	$\sigma_{observed}$ **(mJy)**		σ **(mJy)**
PACS 70 µm	N/A	$q_{phot} = 5.0$	2.26×10^{-3}
		$q_{SDC} = 8.9$	1.42×10^{-2}
PACS 100 µm	0.27	$q_{phot} = 5.0$	1.98×10^{-2}
		$q_{SDC} = 8.7$	1.02×10^{-1}
PACS 160 µm	0.92	$q_{phot} = 5.0$	3.97×10^{-1}
		$q_{SDC} = 7.13$	9.93×10^{-1}

	$\sigma_{observed}$ (mJy)		σ (mJy)
SPIRE 250 µm	5.8	$q_{phot} = 5.0$	2.51
		$q_{SDC} = 5.2$	2.70
SPIRE 350 µm	6.3	$q_{phot} = 5.0$	4.4
		$q_{SDC} = 3.6$	3.52
SPIRE 550 µm	6.8	$q_{phot} = 5.0$	3.69
		$q_{SDC} = 2.5$	3.18

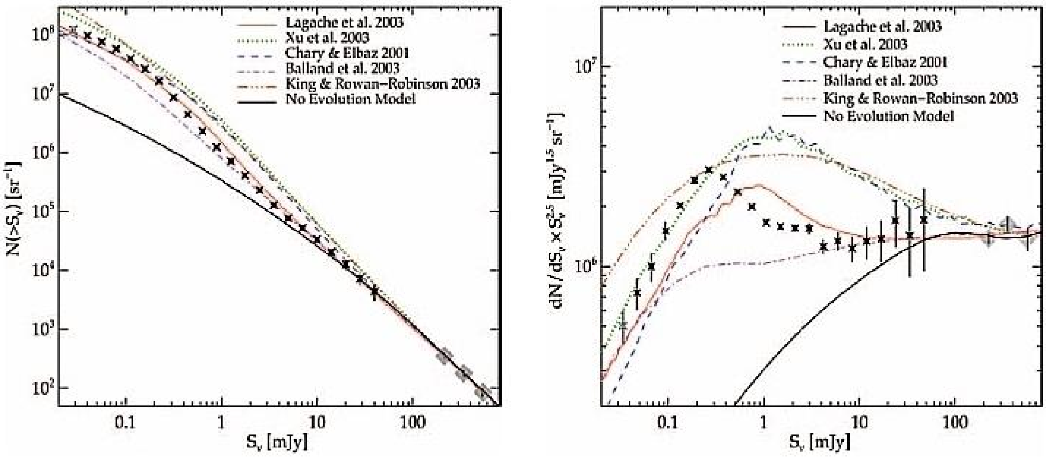

Figure 4.12. Cumulative (left) and differential (right) 24 µm number counts from [RD10]. The differential counts have been normalised to an Euclidean slope, $dN/dS_{\nu} \sim S_{\nu}^{-2.5}$. The curves show predictions from different recent models, including that from Lagache et al. 2003.

4.4. Straylight

The Herschel design was carried out including the instrument optical layout. This approach allowed the level of straylight that originates from the various sources at detector level to be provided directly. Therefore, the straylight requirements were given directly as the straylight reaching the detector. The following apply over the full operational wavelength range:

- **Scattered light from sources outside the telescope FoV**: Taking into account the worst possible combination of the positions of the Moon and the Earth w.r.t. the line of sight (LoS) of the telescope, the extreme values are:
    - Sun-S/C-Earth angle of 37°
    - Sun-S/C-Moon angle of 47°
    - Sun-S/C-LoS angle of 60°.8 to 119°.2 (in the S/C XZ plane)
    - Maximum roll angle of ±1°

  The straylight will be < 1% of background radiation induced by the self-emission of the telescope.
- **Sources inside the FoV**: over the entire FoV at angular distances ≥ 3 arcmin from the peak of the point-spread-function (PSF), the straylight shall be $< 1 \times 10^{-4}$ of PSF peak irradiance (in addition to level given by diffraction).
- **Self-emission**: The straylight level, received at the defined detector element location of the

PLM/FPU straylight model by self emission (with "cold" stops in front of PACS and SPIRE instrument detectors), excluding the self emission of the telescope reflectors alone (but including any other contributor, notably the M2 hexapod), shall be < 10% of the background induced by self-emission of the telescope reflectors.

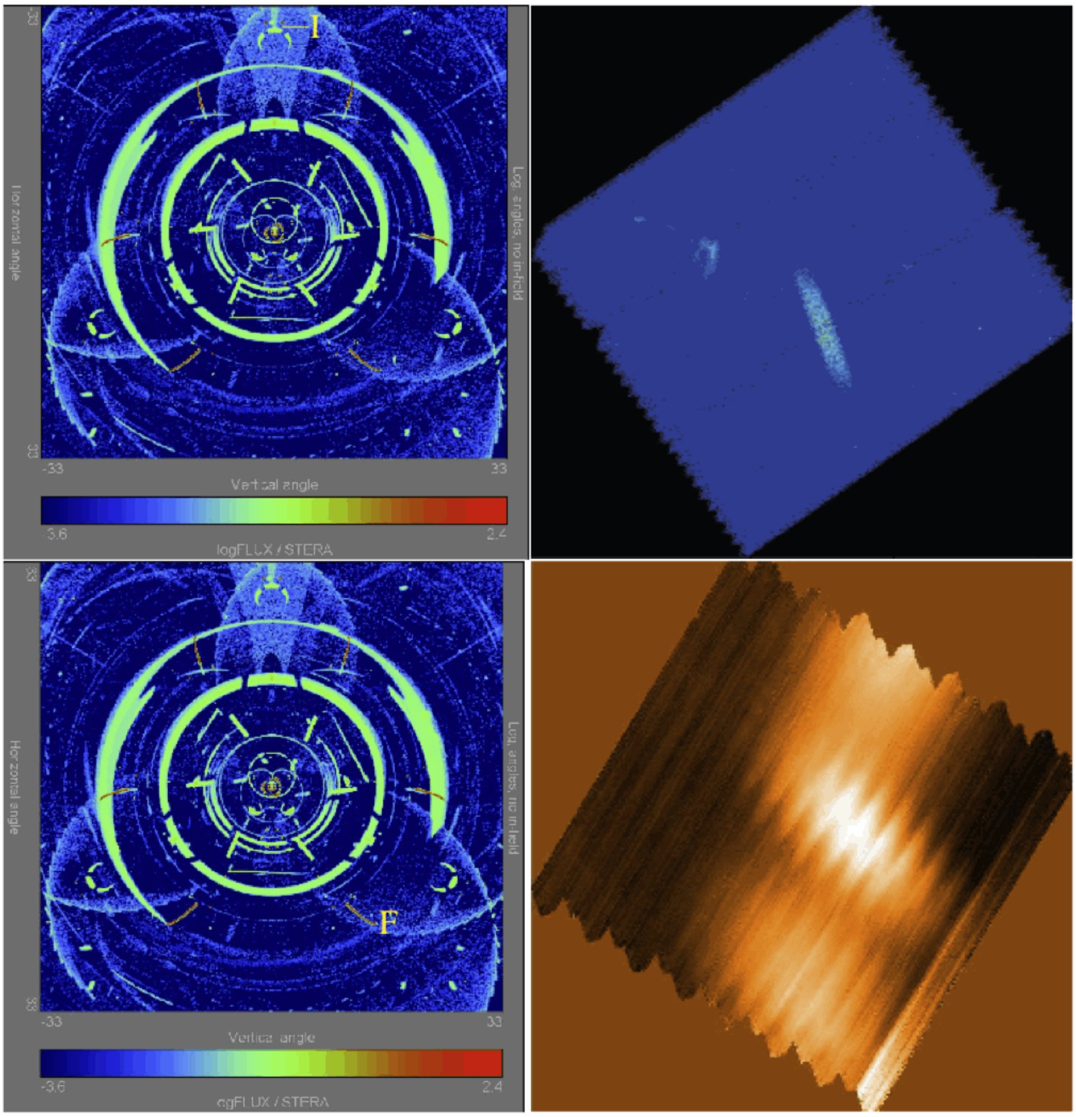

Figure 4.13. Comparison of straylight optical models produced by M. Ferlet (priv. comm.) and observational results. In the top row, a Herschel observation has been planned with Jupiter in position 'I', while in the bottom row the Moon has been placed in position 'F'. In both cases, there is a very good agreement between the model prediction and the straylight results.

According to current straylight analysis for the orbit configuration of Herschel (see [RD3]), for sources outside the FoV, the straylight radiation is within specification, except for small locations on the sky, where radiation reflected from rectangular hexapod structures can enter the instruments directly. These small locations exist primarily for the Moon. Only two minor paths were found which could be applicable also to the Earth. For the worst-case locations of the Moon, the specification is exceeded by a factor 16.4.

For sources inside the FoV, the requirement is met by a wide margin.

Finally, for thermal self-emission, the requirement is not met. Actual values (expressed as a fraction of the background induced by self-emission of the telescope reflectors) are:

- 30% for PACS and 19% for SPIRE (pessimistic case)
- 12% for PACS and 8% for SPIRE (optimistic case)

The HSC created a dedicated working group to study the straylight efffects on the Herschel operations. New models based on the 'as built' optical system were prepared and observations were made to verify the models (see Figure 4.13 for examples of model and observational reality). Improvements in stray light checking in mission planning were made during operations and all cases of possible conflicts were verified manually by the Spacecraft Environment Scientist; where there was found to be a significant danger of degradation of the quality of an observation due to straylight it was re-scheduled for a later epoch. In a few cases late in the mission, when there was no possibility of re-scheduling an observation potentially affected by straylight, it was accepted that such observations had to be scheduled on a "now or never" basis, with a date for execution being choosen that minimised the effects of straylight, provided that the observations was of a point source and that its scientific quality would not be significantly degraded.

Chapter 5. Ground Segment

5.1. Ground Segment Overview

Herschel Space Observatory operations have been conducted in a decentralised manner. As can be seen in Figure 5.1, the Ground Segment has comprised of the following elements:

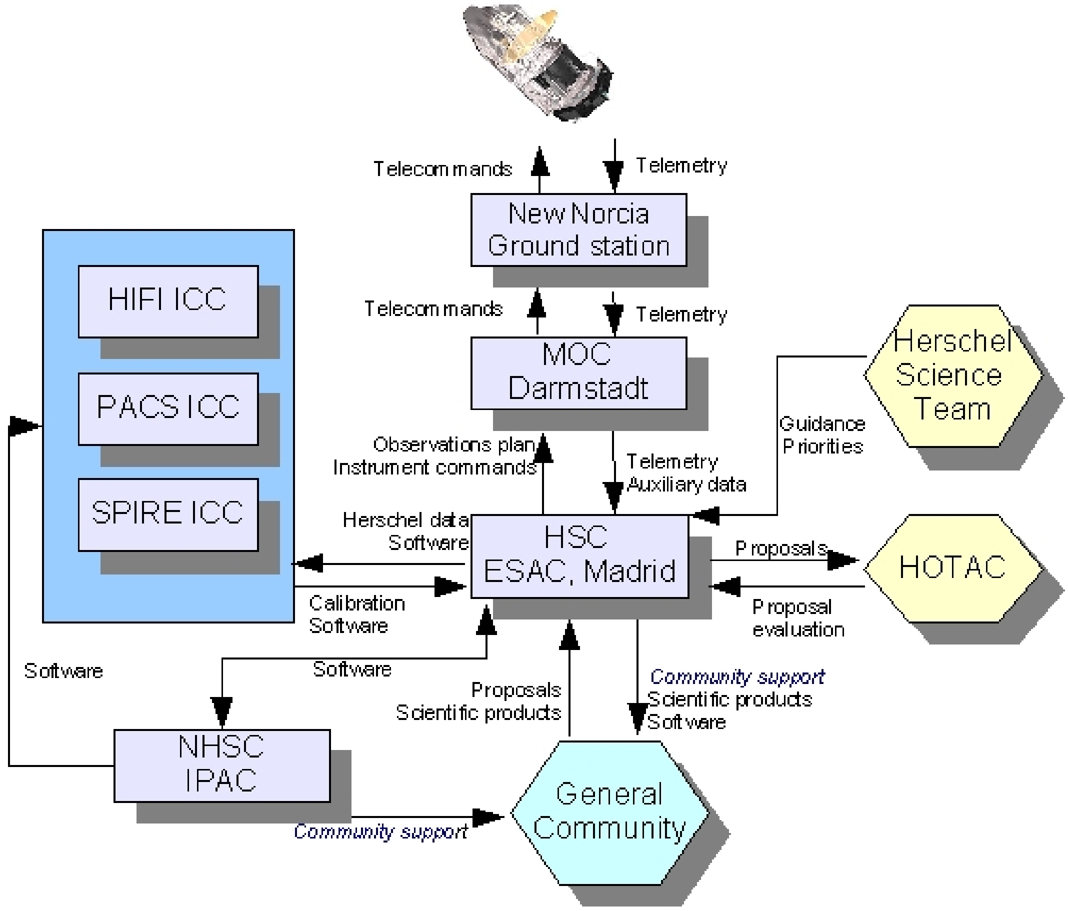

Figure 5.1. Herschel Space Observatory Ground Segment during flight operations. For post-Operations, the MOC branch is no longer used, nor is HOTAC. Otherwise the interactions remain essentially identical in post-Operations to the Operations model, although with modified aims and priorities.

- A Herschel Science Centre (HSC), provided by ESA, located at ESAC, Madrid. The HSC, supported by the NASA Herschel Science Center (NHSC), located at IPAC, acts as the point of interface to the science community and the outside world in general.

  Note
  Note the different convention of spelling of "Centre": this is not a typo. ESA and Herschel use British English spellings as standard, NHSC uses American spellings.

  The HSC is supported by the Herschel Science Team, for the maximisation of the scientific return of the mission, and was supported in Operations, by the Herschel Observing Time Allocation Committee (HOTAC) for the selection of observing proposals.

- Three dedicated Instrument Control Centres (ICCs), one for each instrument, provided by the respective PI. Each ICC has been responsible for enabling the operation of and continues to support the calibration of its instrument in post-Operations.

- A Mission Operations Centre (MOC), provided by ESA, located at ESOC, Darmstadt, which has been responsible for the execution of all in-orbit operations. MOC's role has ended with the end of in-orbit operations, although some support has continued during a transition phase of three months after switch-off, to ensure the final consolidation of data and the successful closure of outstanding issues

5.2. From proposal to observations and exploitation of the data archive

5.2.1. ICC and HSC user support

5.2.1.1. Routine mission support

From approximately three years before launch, when the first call for proposals was made, to the end of helium, the Herschel Science Centre, in collaboration with the ICCs, has provided the information required for the submission of proposals in the Herschel Space Observatory Web site (http://herschel.esac.esa.int). Its role has now switched in post-Operations to one of supporting efficient exploitation of the data archive.

Astronomers are requested to register to access observatory services, which included the capability to submit proposals and continue to provide access to the Helpdesk for the resolution of all issues related to data to Operations and for the retrieval and exploitation of observational data from the Herschel Science Archive. Only registered users can submit Helpdesk tickets, can be co-users of executed proposals and can retrieve data from the archive.

5.2.1.2. Winding-up of post-Operations activities

Support to users from the Herschel Science Centre will continue through to the end of the Archive Consolidation Phase, at the end of 2017, at which point it will become mainly the ongoing work of the ESA archives group, the Satellite Archives Team (SAT). The levels of support from the individual ICCs will drop to a low levels by 2015 and individual ICCs will cease to operate during 2016, with PACS being the first to wind up on March 31st 2016 and HIFI the last, on December 31st 2016. Formal support for the United States community will end with the winding-up of NHSC on September 30th 2017.

5.2.1.3. Herschel Helpdesk

The standard user interface to the Herschel Science Centre has been and will remain the Helpdesk (http://herschel.esac.esa.int/esupport/). As of the end of November 2013, 6626 questions have been dealt with through the HSC Helpdesk. However, as the average Helpdesk ticket leads to approximately five interactions (notifications, enquiry, reply to user, response from user, reply to response, etc) there have been more than 32 000 interactions with users in total through the Helpdesk. The system went live with the first question to the Herschel Helpdesk submitted on Monday 05 Feb 2007 at 06:54 PM (CET).

Helpdesk is a web-based system. This has numerous advantages. As tickets can only be generated by registered users through a web interface, the problem of wasted time and resources in filtering the genuine enquiries from users from spam messages is eliminated. It also makes it easy to identify the Herschel community and to avoid the need for a scattergun approach to user communications. The ticketing system also allows tickets to be assigned a priority on a six-level scale from "low", through "medium" (the standard level for the majority of communications) to "Critical" -- the highest level of urgency, used to flag issues that needed a reply with extreme urgency, for example, scheduling issues that threatened the integrity of observations, or the delivery of schedules to MOC. In post-operations it seems unlikely that any level higher than "Urgent" will be appropriate: this would be for reporting, for example, serious software issues for which an urgent workaround is needed.

Users need to be aware though that, although they receive notification of answers to Helpdesk tickets via e-mail, the return address (hschelp@sciops.esa.int) is a fictious, no-return address: any reply to this address will be undelivered. The user should click on the link in the e-mail and answer the ticket within the Helpdesk system.

A further advantage of the web-based system is that it helps tracking which enquiries are awaiting reply and how long they have been inactive in the system. HSC personnel have taken great pride in minimising the delay in responding to queries as much as possible, but a small fraction of tickets can be "lost in the cracks" -- it was possible to track easily through the Helpdesk system which tickets had not be answered on a day to day basis and take the necessary action, where required. The Helpdesk system would flag any urgent ticket for which a response was overdue.

Throughout operations, Helpdesk tickets were almost invariably read and assessed almost immediately almost 24h a day, even on weekends. Urgent tickets would generally receive a human response of some kind within an hour of submission, if only to reassure the user that a human being was at the end of the line and that the ticket had been seen and was being acted on. While genuinely urgent enquiries are less frequent in post-operations, assessment of new tickets is still carried out rapidly and, where necessary (although this process is usually not seen by the user) classified and brought to the attention of the appropriate expert. While complicated technical issues and software problem reports may take some time to answer adequately, particularly for instrument specialists who have, at times in Operations, become heavily overloaded, many routine enquires have been answered almost immediately throughout Operations, even outside of normal working hours; this service has been offered on a best efforts basis, based on the goodwill of HSC personnel and, only at critical junctures such as the closure of Calls, has it been formalised and then only for a maximum of 48 hours before closure. Even so, even a considerable fraction of routine enquiries have been dealt with in real time, or near real time, for around eighteen hours per day, seven days a week through operations.

Using Helpdesk for user interactions leaves a paper trail that allows responses to be tracked. This enables consistent responses to be given to common queries. However, tickets are not searchable from the Internet (i.e. you cannot Google a question and see a Herschel Helpdesk answer), so they remain private and confidential between the user and the HSC and can only be seen by HSC personnel with Helpdesk access rights.

During post-Operations the intensity of Helpdesk traffic has, logically, reduced considerably however, the typical level of traffic is still 30-40 interactions with users per week.

As post-Operations continue, the HSC workforce will be reduced progressively, with team members being increasingly shared with other missions. However, a full Helpdesk service will continue until the end of post-Operations, although there will be a logical reduction in the level of specialist instrument support, particularly from the start of 2015.

Herschel's Helpdesk will continue to function even after the end of post-Operations, but this will be on a best efforts basis, without specific staff assigned to answer queries.

5.2.2. Proposal preparation and submission

5.2.2.1. Proposal submission and retrieval

Proposal preparation and submission was done through the HSpot tool (see Section 6.2), the Herschel Observation Planning Software. A valid Herschel scientific proposal contained at least one AOR, or Astronomical Observation Request. Each AOR is based on an AOT, or Astronomical Observation Template, which is a pre-defined observing mode, characterised by an instrument configuration and way of operation that have been optimised for the execution of a particular type of observation (see Chapter 6). An AOR was generated when the proposer provided the parameters required for the selected AOT to personalise it to his or her individual requirements, and is equivalent to the term "observation" used in this document.

A proposal submitted through HSpot is stored in the Herschel Space Observatory database. The proposer, and co-proposers selected by the principal investigator to be co-users, were allowed to retrieve, modify and upload their proposal(s) until the closing date of the Announcement of Opportunity. At that time, the database was closed to HSpot, and the HSC distributed the stored proposals to the HOTAC panels. Proposers could check the status of their proposal(s) in relation to the HOTAC review in the Proposal status Web page (http://herschel.esac.esa.int). During the review process, the HSC provided support to the HOTAC and, on request, assessed the technical feasibility of the observations. In addition, a systematic technical feasibility assessment is carried out on all accepted proposals to identify conflicts such as duplications of individual observations or complete proposals with existing or newly submitted proposals (if necessary, HOTAC would consider two or more similar proposals together), technical problems with proposals (observations that were badly, or not optimally designed, or even proposals that were not feasible technically).

All accepted proposals and their associated AORs are available through HSpot and may be downloaded and examined (but not modified and re-submitted) by any registered user. This was done initially to allow new users to design their own observations based on best practices used by the Guaranteed Time users who were generally instrument experts. During the post-Operations phase it may be useful to go back and examine how an observations retrieved from the Archive was designed and executed. The observation can be retrieved in HSpot and a user can set the specific date and time of execution to see exactly how the observation was performed by using the "AOR overlay" option and then animating the AOR.

5.2.2.2. Proposal review and updating

The period of proposal submission before the HOTAC review was called Phase-1. After the HOTAC review the results were made public once collated, checked and ratified by the ESA Director of Science. At this point proposal submission Phase-2 started. In this period, observers were expected to refine their accepted proposals, modifying them following the HOTAC guidelines (this included making any changes that were required after the technical revision of the AORs), remove duplications and prepare the set of AORs that would be executed with the observatory.

Proposals would not be released for execution until the necessary revisions had been made and validated successfully. This meant that the updated AOTs and the latest available observatory knowledge would be used for the observations that would be carried out. Please see the "Herschel Space Observatory Call for Proposals: Policies and Procedures" document for a definition of proposal submission Phase-1 and Phase-2, and for the policies on proposal modifications. The end of proposal submission Phase-2 resulted in a consolidated database of accepted proposals and its corresponding AORs.

5.2.2.3. Late changes to proposals

It was sometimes necessary for a PI to modify the AORs even after Phase 2 has ended. Possible reasons include correcting or optimising coordinates, changes of observing strategy as a result of analysing previous observations, forced changes due to instrumental issues that may have emerged (from time to time the HSC contacted PIs to warn them of instrumental issues that had been detected and to recommend changes of strategy) and the need to replace targets when the original target is found to be unsuitable in the light of new information.

Minor changes could be made by the PI at any time, although to guarantee that the revised AORs could be scheduled they were required at least 4 weeks and preferably 6 weeks before they are due to be executed to allow time for processing (all new AORs go through technical checks) and to allow for the lead time needed when advance drafts of the schedule are prepared. However, more substantial modifications (e.g. changes of target, or changes of observing mode) had to be approved by the Project Scientist before the AORs could be released for scheduling and, of necessity, had to be properly justified, with the exact modifications detailed; these changes had to be coherent with the original aims of the proposal as approved by HOTAC.

5.2.2.4. The post-Operations database

Logically, once EoH was reached, no more science AORs could be executed. It thus made no sense to retrieve and update proposals. Some engineering and a few calibration AORs were executed during the six weeks of post-cryo Operations before passivisation, although these are of no use to general users. Once passivisation had been completed and some clean-up of the Archive carried out (the elimination of observations from the database that were scheduled by the HSC and delivered to MOC, but never executed because of EoH) the database enters its final state and becomes frozen. Users can no longer modify proposals and AORs, but HSC personnel can still make some modifications in the database: this may still be necessary if, for example, subtle problems are identified in a particular observation or observations, necesitating that they be declared failed after detailed Quality Control follow-up.

5.2.3. Data retrieval

After execution, each Herschel observation has had its proprietary period, initially one year for Key Programmes, for which data reduction was more challenging as the quantities of data were greater and the knowledge of the instruments and their calibration correspondingly less. This proprietary period was later reduced to six months. Some observations approved as DDT/ToO have had a shorter proprietary period, or even none at all, if it was deemed that the observations were being performed as a public service and should be made available to any interested user. After the proprietary period has ended, the observations become public. The last data to become public were those from the end of the mission, whose proprietary period ended at the end of October 2013, six months after the end of helium. From the start of November 2013 all Herschel science observations have been public and can be consulted and accessed by all registered users.

While proprietary, only the PI of the data and designated co-users can see browse products or retrieve the data. For proprietary data, retrieval is blocked and the browse products are hidden in the archive to users who do not have PI rights, or who have not been designated as co-Users of the data by the the PI: other users only see an icon warning that the data is not yet viewable. When public, any registered Herschel user can retrieve and publish the data and any person can see the browse products, where available, through the Internet.

Some additional data will be made public in the future. Many calibration observations (see Section 5.3) were taken by instrument teams in standard science modes and can be used for science. The use of standard science modes to take calibration data started even before the first scheduled science observations were made. Definitive lists of releasable calibration observation for each instrument are being prepared as of late 2013. These data will be made available in the archive when the final contents of the release are agreed by all parties.

5.3. Calibration observations

The calibration and cross-calibration of the Herschel instruments is the responsibility of the observatory, in particular of the ICCs and the HSC. The pointing calibration is the responsibility of the HSC and the MOC. Therefore, the preparation and scheduling of calibration observations has been an exclusive duty of these groups. The calibration data required for the reduction and analysis of the Herschel observations is provided to the astronomer in the form of products in the Herschel Science Archive, and is integrated in the Data Processing software.

Calibration and engineering observations were the main components of the schedule during the Commissioning and Performance Verification phases. Their aim was to achieve the necessary understanding of the instruments and spacecraft, and attain the required calibration and pointing accuracies to ensure a proper execution and data reduction of the science observations during the Science Demonstration and Routine phases. Over the duration of the routine phase, up to 15% of the available observatory time has been used for calibration, with the amount of calibration time decreasing towards the end of the mission, as instrument characterisation was completed and only routine trend analysis and health checks became essential. Calibration observations may be based on non-AOT observing modes defined by the instrument specialists at the ICCs and HSC, but in general they have been defined using the AOTs available to the community for science observations.

Calibration observations are in principle public and have been made available throughout the mission to PIs, if justified, sometimes by "buying" the observation against time awarded by HOTAC, which transferred it to the PI along with the corresponding proprietary rights, often "free of charge", by making the observation public and available to all users. However, if a calibration observation is a duplicate of a scientific observation (see the "Herschel Space Observatory Call for Proposals: Policies and Procedures" document for a definition of "duplication"), the corresponding proprietary rights have been applied.

Note

Neither calibration nor engineering observations appeared in the listing of observations that were scheduled for observation and delivered to MOC, but not yet executed (http://herschel.esac.esa.int/observing/ScheduleReport.html), or the observing log of executed observations (http://herschel.esac.esa.int/observing/LogReport.html). The observing log constitutes the official record of executed science observations during the Herschel mission.

Chapter 6. Observing with Herschel

6.1. General introduction

6.1.1. Background

Herschel has been an observatory mission. Thus, as in ground-based telescopes, the astronomer who has requested the observations has had provide all the information necessary to carry them out. These instructions are known as an "Astronomical Observation Request" (AOR), which is made using a standard Astronomical Observing Template (AOT) (see Section 6.4). This information was then converted into spacecraft and instrument commands that were uplinked to the spacecraft to execute the observations. An additional complication with Herschel was that communication with the satellite was normally limited to 3 hours each day, so that all the commands to carry out observations had to be uplinked at least 48 hours in advance of the observations being carried out and had to be executed autonomously. This means that far more detail had to be defined by the observer than for observations in a normal groundbased telescope where details such as the exposure time can be modified at any time. The system was designed to make the highly complex process of defining observations as simple as possible for the observer. The following section describes this process.

6.1.2. Sky coverage over the mission

Herschel was not planned as a survey mission, thus its sky coverage is neither complete, nor homogenous. Herschel observed the areas of the sky that were requested by its users and approved for observation by HOTAC, with no effort to fill in the gaps. Some areas have been covered to great depth, others have rather shallow coverage as a function of the observations that were executed and their aims. Similarly, there was a mix of pointed and mapping observations.

Overall, 9.7% of the sky was covered by Herschel to varying depths with 37 000 science observations taken over 1446 ODs from launch to EoH. You can play and download a 48-second video prepared by Pedro Gomez of the HSC, entitled "Mission Incredible", showing how the 37 000 observations were taken and the sky areas covered, by clicking on this link: http://spaceinvideos.esa.int/Videos/2013/11/Herschel_s_37_000_science_observations. Time -- measured in both UT and ODs -- is shown at the top of the screen, while the labels of all the AORs to be scheduled scroll across the bottom of the screen and the planets move along the ecliptic.

Figure 6.1. The final state of Herschel sky coverage. All 37 000 successful Herschel science observations are represented with their correct sizes. It is immediately obvious that, although the AOR centres are re-

markably homogenously distributed around the sky (see Figure 3.8), in terms of actual sky area covered, the sky coverage is quite variable, with certain areas such as the Galactic Plane and Galactic Poles particularly well covered.

The final state of sky coverage is shown in Figure 6.1, with all 37 000 science observations represented. Although the AORs centres are remarkably homogeneously distributed on the sky (see Figure 3.8), in terms of sky area covered, the Galactic Plane, Ecliptic and Galactic Poles are all particularly well covered. There are also a few areas of the sky where there are very few observations, such as the constellation of Lacerta or the large hole in Auriga/Lynx.

6.2. Introduction to HSpot

The astronomer's interface with Herschel was an observation planning program called HSpot, or HerschelSpot. HSpot allows the astronomer to define targets and observations, to calculate the time required and likely s/n and to submit a proposal with the requested observations, as well as offering many other useful functionalities to plot observations on the sky, see how they were executed, check an observing field against images and catalogues at various ranges and to look for other observations taken of the target. At any stage of this process the work in progress can be saved and recovered later.

HSpot has many useful funcionalities for exploring how observations were carried out and overlaying catalogues over images, or AORs, or explorimng sky coverage in a region of sky, among many other possibilities. This gives it a a big legacy value. in post-Operations HSpot is only maintained on a minimal, best-efforts basis, but the HSpot 7.0 version of the program will our legacy version.

HSpot has been adapted from the original Spitzer Space Observatory SPOT program and thus will be familiar to Spitzer users. The part of HSpot directly adapted from SPOT is known as the "Spot Core" of the program and is maintained by IPAC (about 80%), while the HSC has maintained the layer of Herschel-specific functionality (about 20%); HSpot incorporates a total of more than 30 man-years of work between the two centres.

HSpot can be downloaded from the Herschel Science Centre web page at the url:

ftp://ftp.sciops.esa.int/pub/hspot/HSpot_download.html

Alternatively, select the "Tools" option from the left hand menu of the Herschel Science Centre web page and "HSpot download" in the tools page.

HSpot is eminently user-friendly, simple to use and has many functionalities that are of interest even to non-infrared astronomers. New users can generally familiarise themselves with the main functions in an hour or so of simply playing with the program.

6.2.1. Keeping HSpot up to date

6.2.1.1. Operational versions of HSpot

HSpot has been updated regularly during Operations. For the OT2 Guaranteed Time Call in 2011, a completely new and revised version was released (HSpot 5.3), including literally dozens of minor and major updates since the previous Open Time Call and also numerous updates to the underlying Spot Core. Version 6.0 was then released for the OT2 Open Time Call; this version was updated for Phase 2 of OT2 as HSpot 6.1 and was further updated during 2012 as HSpot 6.2. A final operational release of HSpot 6.3 was made for the final few months of observing. The final release of HSpot that is planned is the 7.0 version, in autumn 2013, for which the only actual changes are a change in the underlying software version on which HSpot is built and an update in the On-Line documentation for post-Operations.

A planned update of the installation software to allow HSpot to install successfully on Windows 8 computers has had to be deferred, although Windows 8 is not anyway, at this time, a supported operating System. It is possible that a Windows 8 fix, which will also address installer issues with Mac Mountain Lion, will be implemented to coincide with the HCSS 12 release planned for March 2014, with a new HSpot release, built on HCSS 12, being made at this time.

Occasionally, unexpected issues came to light, requiring a new update of HSpot between major releases, in which case a new release was made. Not all releases were public releases for the astronomical community; many contained only expert user updates of interest to HSC staff and calibration scientists at the ICCs, hence the jumps in numbering of versions that users may have noticed between user releases. These incremental changes were included in the regular user releases, although they were not visible to the majority of users. The default setting in HSpot is that it will download these updates automatically and offer them to users when a new user release is detected. It was strongly recommended to users not to change this option, as it could have led to submitting or revising their AORs against a wrong HSpot version, or to having incorrect time estimates for them due to using an outdated time estimator (in extreme cases it could have even led to a proposal being rejected automatically by the Proposal Handling servers at HSC).

Each time that you open HSpot, it will connect to the HSC server and check to see if a new version is available. If one is found, you will be offered a choice of closing HSpot immediately and re-opening it with the new version, of waiting to install the new version, or of refusing the update (in which case automatic updates are disabled in the future). You are strongly advised to accept the update immediately; normally it will be installed and operational in under a minute. It is still possible that, even after the HSpot 7.0 release, small updates will be needed very occasionally to maintain critical funcionality.

Similarly, occasionally a new time estimator version may have been announced on opening HSpot; the time estimator links the AORs to the latest instrument control software that set the parameters for each observation. Normally, by the second half of the mission, time estimator changes did not affect the duration of observations, but it did effect essential parameters in the set-up of observations. When the time estimator version was updated, the time estimate for your previously prepared AORs would be shown in red; it was essential to submit all proposals against the latest time estimator version - to do this, it was only necessary to re-run time estimation before submission so that all the time estimates were shown in black font. When loading old AORs from disk their time estimation may be seen as out of date.

Executed AORs downloaded from the HSC will show the time estimate resulting from the time estimator version that the AOR was executed with, which may be a particularly old version; on occasion you may find that this causes significant changes in time estimates when these AORs are updated with the final time estimator: when a proposal was updated and re-submitted, the system ignored all changes to AORs that had already been scheduled, including timing changes, hence on downloading an updated proposal, observations executed with a previous time estimator version would show, correctly, as out of date. This means that all AORs in the database are effectively frozen with the configuration that they had at execution, meaning that they can be retrieved and examined to see how they were executed by the observatory.

6.2.1.2. The legacy version of HSpot

In post-Operations there is no need to submit and modify proposals and HSpot will only be used for legacy purposes. The legacy version of HSpot is 7.0 and includes a raft of final, small updates. It is not expected that there will be a further release (although a minor new release may be made in Spring 2014 to build HSpot 7 on the new HCSS 12 software, see Section 6.2.1.1), but it cannot be discounted that some small changes may be required to keep HSpot operating in the future.

The biggest potential difficulty for HSpot is future changes in operating systems and in Java creating incompatibilities. Plans are in place to virtualise HSpot so that it will remain operational even under future operating systems and Java versions.

6.2.2. Will HSpot run on my computer?

HSpot has been developed to run on the three main operating systems currently in use: Unix/Linux, Windows and Mac. The development work was carried out on Solaris and ported to these operating systems and the system has been extensively tested and used operationally by many hundreds of users for some seven years. We thus know that HSpot should run reliably on all the principal operating systems available to users. For each operating system certain common platforms are supported. Users are strongly urged to use these standard combinations of operating system and platform, as no guarantee can be offered that HSpot will run correctly on other combinations and no guarantee can be made of support for other platforms; occasional serious issues have been reported with other platforms, particularly for Linux users due to the wide variety of platforms available for Linux (it is physically impossible to test or to support all possible Linux platforms), with new platforms appearing late in the mission that occasionally had serious compatibility issues. Similarly, users will understand that, for example, the Windows version of HSpot has been extensively tested on Windows XP, Vista and Windows 7, but was developed with Windows XP.

Detailed information on the operating systems and platforms supported can be found in the HSpot manual. HSpot runs under Java and users are strongly advised to ensure that all updates and patches of their operating system are installed.

6.2.2.1. Java issues

HSpot will only run on Java 1.6 and its later updates. Older Mac machines that do not support a Dual Core and 64-bit architecture will not install HSpot as they cannot run Java 1.6. Due to the way that Mac handles Java, Mac users have occasionally experienced minor problems with HSpot that Windows, Solaris and Linux users have not.

As there is already a Java 1.7 release, some issues may arise when trying to use HSpot on a computer with Java 1.7 installed. On Windows and Linux HSpot is packaged with the appropriate Java 1.6 release, so HSpot should run correctly. Java is *not* packaged with HSpot for Mac and uses the Java version packaged with the operating system. Once Mac launches its successor for Mountain Lion, it will no longer be possible to have Java 1.6 installed on a Mac and thus, no guarantees can be made that HSpot will still run on Mac.

6.2.2.2. Windows 7, 8 and later releases

HSpot is tested on all Windows releases up to Windows 7. We know that Windows 6.3.2 -- the final Operations release -- will not install on Windows 8 machines. The reason is that the version of the InstallAnywhere software used to install HSpot with which this release was built was not Windows 8 compatible. Technically, the jre (Java Runtime Environment) will not install correctly with this version of InstallAnywhere. Although a newer version of InstallAnywhere is available, other issues have been found with it so, for the time being we are unable to support Windows 8. Other solutions are being investigated actively for these issues.

6.2.2.3. Other OS releases

We are aware that the InstallAnwhere software will not work with Mac Mountain Lion and that the supplier has withdrawn support for this OS. As new platforms and operating systems become available such issues may be become more commonplace. Again, we are working actively on finding a more permanent solution.

6.2.3. Proposal presentation and retrieval

Proposal presentation has been extremely simple with HSpot. Once the observations to be carried out were defined and saved, the proposal could be submitted quickly and easily from the "Tools" menu. A submitted proposal could be retrieved before the deadline for submission and revised as many times as required; this allowed users to submit a draft and then update it continuously so that, even in case of disaster (a local hard disk failure, the Internet falling over just before the submission deadline, etc), HSC always had a valid latest version of the proposal. Similarly, a ToO/DDT proposal could be revised as many times as necessary before the Project Scientist subjected it to evaluation when informed that the definitive version was available for this evaluation process.

Similarly, accepted proposals are public documents that may be retrieved and examined by any registered user. The original intention behind this was to allow new users of Herschel to be able to study successful proposals and their AORs to be able to prepare a more competitive and technically better proposal in the Open Time Calls. However, it has been found that it is often very useful to be able to look at a proposal to understand the AOR design, samples, s/n calculations, etc. also when studying data in the archive.

AORs and abstracts can be retrieved via the "View Accepted Proposal" option in the "File" menu of HSpot. The core of the proposal (i.e. the scientific justification) can though only be seen by the P.I. of a proposal and his or her designated co-Users.

To submit a proposal, apart from the AORs (that is, the source information, instrumental configuration, exposure time, etc. for each object to be observed) the proposer needed a text file with the proposal abstract (maximum 2000 characters, including spaces), which could be read in directly, a PDF file of the scientific justification (limited to a maximum of 5Mbt and prepared with the latest version of the HerschelFORM PDFLatex package that is available on the Herschel Science Centre webpage) and to give basic information such as the proposal title, list of co-Is and the observing call that the proposal is responding to.

When a proposal was submitted, HSpot would confirm that it had been transmitted correctly and, on completion of processing, an e-mail was received from the HSC Proposal Handling System confirming its successful receipt. The time taken to generate and transmit the acknowledgement e-mail was a strong function of the system load. When the HSC servers were heavily loaded close to a call closure, the acknowledgement e-mail could take tens of minutes or even, in extreme cases, a few hours to arrive. Until this e-mail was received, users were not able to retrieve and update the latest version of their proposal. All proposals that arrived were logged with the time of submission and the HSC knew that a proposal was in the system, even if the formal acknowledgement had not been received. Once formal Calls for Proposals ended, the load on the system was normally extremely low and a proposal would be acknowledged rapidly. However, proposals with large numbers of HIFI observations, especially Spectral Scans, were much slower to process than PACS or SPIRE photometry proposals and could still take some minutes to be processed and acknowledged.

The database is now frozen and cannot be changed by users save for minor details of adding co-users for data access (this only makes sense for proprietary data so, from November 2013, there will be no need to make such changes any longer and this capability will be removed as it makes no further sense to have it), or to allow them download complete proposals. Proposals, logically, can no longer be submitted or updated, but they and their AORs can still be retrieved and examined with HSpot. Retrieving and examining the AORs from a proposal may be useful to understand how data retrieved from the archive was obtained.

6.3. Types of target

HSpot deals with two fundamental types of target: fixed targets and solar system objects.

6.3.1. Fixed targets

A fixed target is any object that does not require a differential tracking rate. This can be a star, a galaxy, an AGN, etc. Herschel works with Equatorial J2000 coordinates and only target entry in Equatorial J2000 will be accepted (this was set to facilitate checks for duplicate pointings [two or more users requesting similar observations of the same object, or region of the sky], which are extremely complicated if many coordinate systems were used for target entry). If the source is known to NED or SIMBAD these coordinates are used, if not, the user must enter a J2000 R.A. and Dec. On some occasions, for nearby stars, the proper motion of the target may become important; this can be entered in HSpot if necessary, once again, the epoch must be in 2000 coordinates. All fields can be edited after name resolution.

Some types of observation could be very unforgiving with the coordinates, particularly PACS spectroscopy and HIFI observations in the high frequency bands (Bands 6 and 7) are sensitive to poor pointing accuracy and may have been affected seriously if the quality of the position of a fixed target was low. Many faint far infrared targets do not have good coordinates and the optical position and far infrared position of a particular target may not necessarily coincide: this should be carefully bourne in mind when comparing Herschel and optical data: a discrepancy may not necessarily indicate poor pointing, it may be simply an indication that the optical and far infrared source do not coincide.

Typical problems found when entering user-defined coordinates for targets were missing "minus" signs in declination and, where several positions are available for a target, picking an inaccurate one.

Warning

Caveat Emptor! The PI was (and is) responsible for the accuracy of the submitted coordinates and for ensuring that the coordinates supplied were good enough to obtain the requested data. Observations that failed because of bad coordinates could only be repeated if the time required came out of the proposal's HOTAC allocation, requiring other observations to be cut in duration, or even removed completely from a proposal.

The pointing performance of Herschel varied considerably through the mission as it was better understood and refined. Detailed information of the pointing performance and its evolution can be found in Section 2.4.5.

6.3.2. Moving targets and their treatment

A moving target is a solar system object that requires a differential tracking rate to be programmed. On target entry the user needed to select the "Moving" tab and resolve the NAIF ID of the target name. The Herschel Mission Planning System used the NAIF ID to calculate coordinates for the time of observation and to calculate the differential tracking rate required. This needed to be less than 10 arcsec/minute at the date of observation (this limited the capability of Herschel to see objects passing very close to the Earth, although faster rates up to and even slightly exceeding 30 arcsec/min could be permitted, on a case-by-case basis, if scientifically justified). User entry of target coordinates is not permitted, as any solar system object with a reliable enough orbit to have been observable by Herschel will have a NAIF ID.

Around 800 moving targets (satellites, comets, asteroids and TNOs) are in the HSpot database. More than a million have been catalogued, but it is obviously impractical to store all of them in HSpot as most were not observable by Herschel. To observe a solar system object that was not in HSpot, for example, a new discovery, it was necessary to send a Helpdesk ticket requesting that it be added. A minimum of two or three working days were needed normally for it to be included and for the ephemeris to be linked to HSpot.

The NAIF ID is shown for solar system objects in the HSpot "position" column rather than an R.A. and Dec. No position data is given for solar system targets in the observing log as this is calculated 'a posteriori" when the data is processed, based on the ephemeris used in downlink (which may not be the same as the uplink ephemeris used to schedule the observation) and the timing data for the observation as executed.

6.3.2.1. What is a NAIF ID?

NAIF is NASA's Navigation and Ancilliary Information Facility. This offers an information system called SPICE for spacecraft navigation. SPICE uses a unique 7 digit identification code for all natural solar system bodies, while spacecraft are identified with a negative integer code. Because of the simplicity for this system of ID codes and given the increasing possibility of confusion of objects (for example, there are both planetary satellites and asteroids named Io, Ganymede and Dione and increasing numbers of asteroids are later found to show cometary activity and may receive multiple designations), it is increasingly used for telescope scheduling. A short summary of information about NAIF IDs is given in the relevant secion of the HSpot Users' Manual on the Standard Ephemeris for moving target entry.

6.3.2.2. Solar system object ephemeris accuracy

When a Solar System Object has a well-controlled orbit of high accuracy (for a periodic comet this means two returns for which a successful linkage has been made, for a asteroid or minor body it usually means observations at a minimum of 6 or 7 oppositions, apart from Earth-crossing objects for which the criterion is typically 3), it will receive a number from the Minor Planet Center. A numbered comet has a designation such as 190P/Name, while an asteroid receives just a number. An unnumbered asteroid has a NAIF ID starting with a 3. Objects with such a designation have a relatively low accuracy ephemeris that may be considerably in error when extrapolated even a short time into the future. As an example, even an object with three oppositions may have a position that has a 3-sigma error of more than 60 arcseconds when extrapolated 5 years into the future. If the spread of observations is unfavourable, or there are few astrometric observations, it may not even be possible to obtain a good ephemeris extrapolation with a 3-opposition orbit. With 4 oppositions the 3-sigma error in the extrapolated position may still be greater than 20 arcseconds over 4 years. This means that faint objects that have not been observed recently may have been difficult to locate and identify with Herschel and thus were high-risk observations. It also meant that an object may not have been centred on the detector, so the resultant data quality may have been defficient.

For comets, non-gravitational forces can make predictions quite inaccurate for returning objects, even when they are well-studied and the closer that they come to Earth, the worse the situation is in terms of ephemeris accuracy. Experience has shown that close approach comets always needed ground-based astrometric support campaigns and late-time ephemeris updates to make observations possible with Herschel and that even updating the ephemeris with the latest data 4 days before the observations were to be executed (this was the latest possible time for update to be made to allow it to be processed in time for Uplink) and re-planning the observations was not always enough to avoid positional errors of several arceseconds between the expected and the actual position at the time of observation if no high-precision radar observations were available to tie-down the position with exactitude. As an example, the difference between the predicted and observed postion of 45P/Honda-Mrkos-Pajdusakova was 7 arcseconds, despite an ephemeris update four days before the observations were executed, as no radar data were available until after the execution of observations with Herschel.

When re-processing data in bulk re-processing at HSC, the best available ephemeris is used. As in the case of 45P/Honda-Mrkos-Pajdusakova, this may be considerably better than the one that was available at the time of observation. Observations in the Archive thus are processed with the best possible positional information at the time of processing, not the best that was available at the time of observation.

6.3.2.3. What accuracy of ephemeris is required?

Three problems may be present when there is uncertainty in the ephemeris. In approximate order of increasing importance these are:

- Possible errors in the required tracking rate.

  In general the tracking errors should be kept below 1 arcsecond during the observation. This could be a problem with long observations on objects moving at high velovity on a strongly curved track for which the interpolation of the position may not be good enough. In the end though, it seems that this has not been an issue for any solar system object observed by Herschel.

- Difficulties with photometry

  For the observation to have been carried out successfully, the target must have been centred in the array to within a certain level of accuracy. If this was not achieved, it may be difficult, or impossible to obtain good quality photometry.

- Problems with target identification

  Not all Solar System Objects have suitably accurate ephemerids to have been observed successfully at all; occasionally there may be errors of tens of seconds of arc, minutes or even, for a few objects, degrees in the ephemeris position. In the HSpot Users' Manual a list of solar system objects included HSpot is given in which flags objects with deficient ephemerides at the time of launch. It was always recommended to check though to see if a better orbit was available when requesting time.

In detail, the issues that users may find when ephemeris information is uncertain are:

- Tracking

  In general this should not be a problem with distant objects, it may become a serious problem with more nearby ones, particularly Near Earth Objects where it may be difficult to keep the target accurately centred.

- Photometry issues

  For PACS photometry, the source position must be known with high enough precision that it should fall within a bolometer matrix of 52x52 arcseconds. In practical terms this means that the following criteria of positional accuracy should be fulfilled.

  -- For aperture photometry: 15 arcseconds.

  -- For PSF fitting: < 10 arcseconds

  For SPIRE the main consideration is that the FWHM of the detectors is 18 arcseconds and the jiggle amplitude 6 arcseconds: if the positional error is greater than the jiggle amplitude there will be light losses.

  For HIFI it should be remembered that the smallest aperture (that of Band 7b) is 13 arcseconds and for PACS spectroscopy the pixel size was of the order of 9 arcseconds, thus necesitating, in both cases, centering at the arcsecond level to avoid light losses.

- Target identification problems

  For numbered asteroids the ephemeris should be of sufficient precision in almost all cases.

  For unnumbered asteroids and minor bodies it has been essential to take astrometry to refine the orbit before observations could be attempted with Herschel.

  For numbered and ToO comets, recent astrometry may have been essential, depending on the case. A numbered comet has almost invariably required post-recovery astrometry to refine the orbit before observation could be attempted and any close-approach object will need extensive late-time astrometry. Recently discovered comets with a short orbital arc have invariably required late pre-Herschel observation astrometry to refine their ephemeris.

6.3.2.4. Potential problems with moving targets

Standard practice at the HSC was to download the ephemerids for Solar System Objects from the JPL Horizons database every 4 weeks, as our planning was done in cycles of 2 weeks. This means that in an extreme case the ephemeris information for an object may have been as much as 2 months out of date. Normally this did not matter, as the errors will be too small to be significant for Main Belt asteroids. If the orbit is well-defined, the error is usually almost entirely in the direction of motion, so a Solar System Object (asteroid or comet) will reach a given point in its orbit slightly advanced or slightly delayed with respect to the ephemeris prediction. For Main Belt Asteroids (MBAs) advances or delays of 10 or more minutes of time in the position along the orbit are not unknown; but, at a typical distance of observation of 2AU from Earth, this translates into a very small error in the actual observed position on the sky.

For objects that come closer to the Earth and move more rapidly, the error in the ephemeris may be much larger, even though the absolute precision of knowledge the object's position may be an order of magnitude better than for an MBA if it had been intensely observed. In the case of Comet 103P/Hartley 2, at the time of observation it was found to be almost 30 arcseconds away from the ephemeris prediction published only 2 months previously. For Near Earth Asteroids and comets an ephemeris may become effectively completely unusable in a week or less. Observers had to be aware that it was their responsibility to ensure that there is sufficient knowledge of the ephemeris of a target for effective scheduling and to warn the HSC sufficiently in advance to take the necessary measures to schedule with the most up-to-date available information.

Normally observations were planned and sent to MOC for uplink to the satellite a minimum of 2 weeks in advance and more usually at least 3 weeks in advance. If knowledge of the object's position was likely to be insufficient at that time, the observer was expected to request -- in advance -- that the observations be re-planned closer to the date of execution, taking advantage of a better ephemeris. Such cases required careful forward planning, both at HSC and MOC to ensure a late-time re-delivery of the observations could be made and processed successfully. On numerous occasions it was also necessary to coordinate carefully the re-delivery with additional ephemeris calculations and updates at JPL, who frequently made exceptional efforts to support the observation of solar system targets with Herschel, including rapid processing of radar observations from Goldstone to make a highly accurate ephemeris available within a few hours of an observation at Goldstone, without which some observations would have been totally impossible to execute successfully.

It was the responsibility of the observer to warn the HSC, via a Helpdesk ticket, of potential ephemeris problems that might have affected scheduling of a Solar System Object. This needed to be done far enough in advance to be taken into account in the standard planning cycle and for the necessary actions and additional deliveries to be discussed and agreed with MOC. However, no re-planning of observations could be contemplated less than 4 days in advance of their execution; if the ephemeris was potentially not robust for 96 hours in advance of execution, the observations had to be designed to be robust enough to compensate for any positional errors that may have occurred.

It was not sufficient to assume that the HSC would automatically spot all potential conflicts with Solar System Objects in advance, although this was part of technical assessment of proposals and every effort was made to anticipate such conflicts.

6.4. AOT entry

6.4.1. Using AOTs

An AOT is an "Astronomical Observation Template". This will be familiar to users of ISO and Spitzer. An AOT is a standard observing mode with an instrument that can be translated into instructions for the spacecraft to carry out the observations autonomously. Herschel observed autonomously between DTCPs, so each observation had to be carried out in a standard way that the spacecraft could understand. Thus, for each of the instruments only pre-defined types of observations could be carried out. The astronomer produced an AOR (Astronomical Observing Request) by taking an AOT and customising it for the required observations.

Following the experience of ISO, the number of AOTs was deliberately restricted to allow observers as many options as possible, without requiring an unwieldy number of observing modes to be validated and calibrated.

The first stage in AOR entry was to define the target. If it is a known object, its name can be resolved with SIMBAD or with NED or, for a solar system target, as a NAIF ID. For unknown names (e.g. start points for scans), J2000 coordinates must be supplied by the observer. After defining the object, the observer needed to check that it was observable by Herschel by calculating its visibility windows. It was also necessary to bear in mind that when defining an observation further, the observer could end up limiting its visibility to part of the target's unrestricted observing window or, worse, making it completely unobservable by putting constraints that were too strong on the observation.

Once the target is defined the observer must then select the required instrument and AOT to be used. Nine basic observing modes were supported: for HIFI, single point (point source spectrophotometry), mapping and spectral scans; for PACS, photometry, line spectroscopy and range spectroscopy; for SPIRE, SPIRE photometer and spectrometer; and the SPIRE PACS Parallel Mode (this last mode was only available for fixed targets and could not be used for Solar System Objects). Each of these modes is further subdivided, HIFI, for example, offered a choice of fourteen different mixer bands. PACS photometry allowed three variants including the mini-scan maps which replaced point-source photometry early in the mission and chopped raster maps. SPIRE Spectrometer offered point source and raster maps, three choices of image sampling, and four choices of spectral resolution, etc. HSpot will guide you through this process of definition with a series of pull-down menus and pop-up windows. All these options still work correctly in post-Operations and can be used to investigate the effects of taking data in slightly different ways to understand your observations.

Warning

Although Point Source photometry may be selected through HSpot, it was not offered as a standard science mode in routine operations and remained only for specialist calibration applications, save for a very few science observations executed in this mode early the mission, which served to show that its sensitivity was lower than had been hoped. Point source photometry for science data was carried out using mini-scan maps only, which are far more sensitive for the same integration time, although calibration observations taken in Point Source mode are fully calibrated and can be used for science.

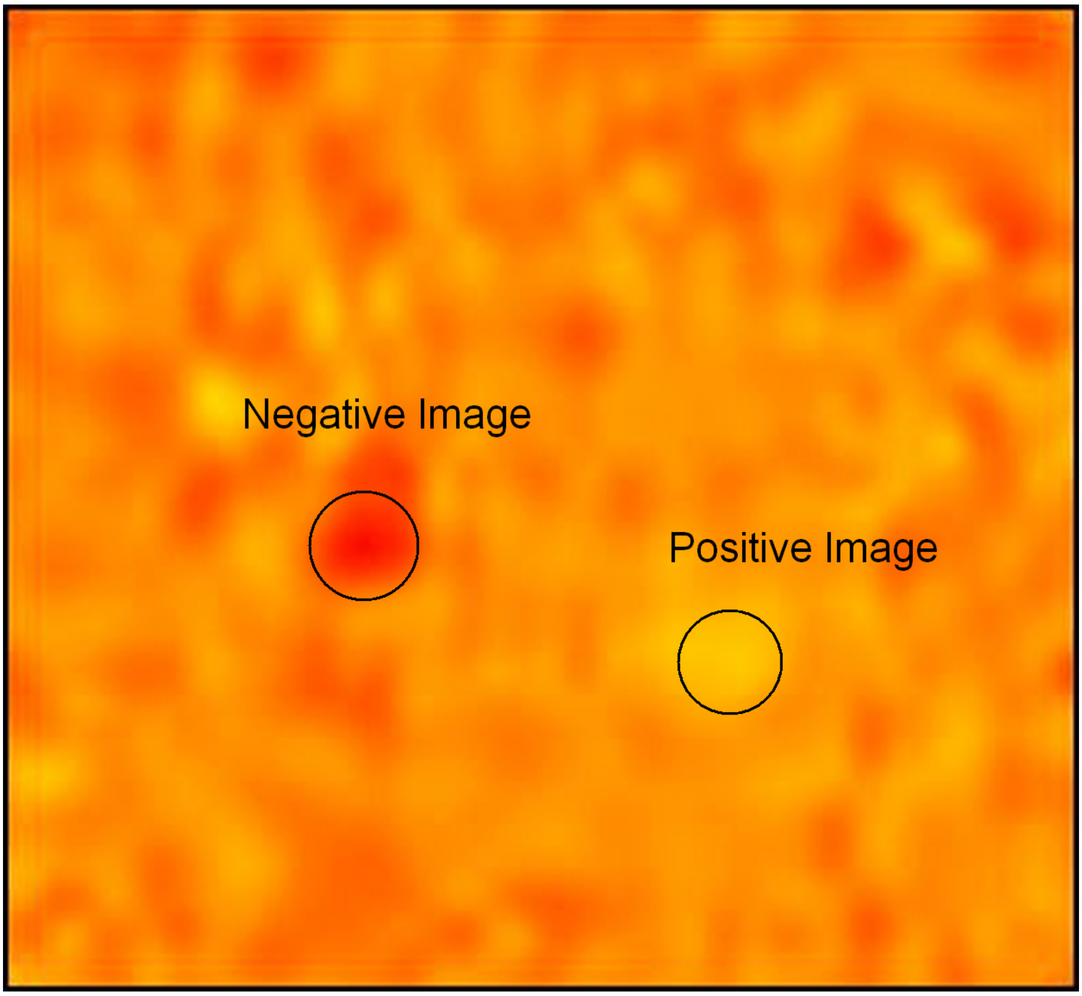

Figure 6.2. Beating the confusion limit for Solar System Objects. This SPIRE image of Makemake uses the technique of subtracting the background using the shift in position between two epochs and subtracting one frame from the other. The trans-Neptunian Object (TNO) Makemake has an estimated flux of 15mJy at 250 microns, with a confusion noise sigma of 6mJy. A weak, but clear detection is obtained in this image with 15 repetitions (NB: the confusion noise is effectively reached in 2 repetitions). This is thought to be the faintest target to be detected with SPIRE.

For each observation there was a basic minimum unit of observing time required; the observer needed only specify how many repetitions of this unit time are required -- obviously greater sensitivity is obtained through more repetitions (four integrations will give twice the sensitivity of a single one -- although for SPIRE, once the confusion limit was reached, you could not attain better sensitivity however long you integrate, although it was possible to obtain better signal to noise), but the observation took longer for a relatively small gain. At any time the "Observation Est..." (Observation Estimate) button can be pressed and HSpot will give an estimate of the total time that the observation would have taken, including the overheads involved, with a break-down of information about the observation. If the total length of the observation exceeds the maximum permitted, HSpot will give a warning that the observation duration is out of limits.

The one exception that allowed the confusion noise to be beaten is for moving targets -- Solar System Objects (SSOs). By using an off position we can subtract out the background and thus eliminate almost completely the confusion noise (see: Figure 6.2). For a sufficiently long exposure of an SSO, the target will move sufficiently during the exposure that the end of the exposure can act as the "off" for the start; this has proved very effective for detecting extremely faint SSOs down to the limiting sensitivity of the telescope.

The observer could vary the parameters of the observation (more or fewer repetitions, nodding on or off, larger or smaller chopper throw, a wider or narrower range of wavelengths or length of scan, etc.) and see how the time estimate varied. Once an acceptable combination of parameters had been found the observer accepted the parameters that were defined to fix the AOR; this AOR can however be modified later, if necessary.

When a proposal was submitted, HSpot took the currently defined list of AORs and linked them to the proposal. It was thus essential to ensure that the correct AOTs and AORs were defined and that the source visibility and observing time were correct for each target.

6.4.2. Full and limited visibility in HSpot

The star tracker (the telescope's autoguider) was pointing in the opposite direction in the sky to the telescope. So, when an object lay in the hemisphere away from the Sun the star tracker was pointed into the sunward hemisphere. It was found that when the solar elongation of a target is greater than 105 degrees, sunlight illuminated the bottom of the S/C platform, leading to heating and thermal distortion of its structure. This distortion caused progressive guiding errors that affected not just the current observation, which would see progressive drift in the tracking, but also all the following observations until the the S/C structure had stabilised. On return to cold attitudes, the process was reversed, with the pointing slowly returning to nominal, again causing progressive pointing drift, thus a long observation at hot attitudes would affect not only the observation in question with a progressive pointing drift, but also later observations.

After in-flight study it became obvious that data quality would be compromised by long exposures in the region of solar elongation from 110 to 119.2 degrees (the "hot zone"). HSpot showed observations in this range as having "limited visibility". In practice this meant that according to the degree of incursion into the hot zone increasing strict limits would be put on scheduling. As a rule of thumb, observations longer than 1 hour would not be scheduled in the hot zone save in exceptional cases.

Observations in the hot zone were scheduled at the end of the observing day, unless that were very short incursions that went only slightly into the hot zone of it could be demonstrated after careful assessment by the Spacecraft Environment Scientist and by the Herschel Mission Planners that they would have no significant impact on the quality of observations that followed: this allowed the star tracker base plate to cool during the Daily Telecommunications Period so, only, in normal circumstances, only short and extremely urgent observations will be scheduled in this hot zone. Time-critical observations longer than an hour that only impinge slightly into the hot zone (e.g. SSOs) were treated on a case-by-case basis.

With the improvements in Herschel's pointing performance, it became obvious that degradation of pointing occurred at smaller solar aspect angles than previously thought, being particularly important for PACS spectroscopy taken after observing in a "hot" region of the sky, which could be seriously compromised by even small errors in pointing. This obliged us to reduce the area of the sky considered to be "cold" and thus available for unlimited scheduling, considering any observation at a solar elongation greater than 105 degrees to be warm (limited to short observations that would not affect later pointings) and greater than 110 degrees as hot (only to be scheduled if scientifically critical and, only then, taking all necessary scheduling precautions, including scheduling at the end of the day and placing only pointing-robust observations immediately afterwards). From late July 2012 HSpot showed 75% of the visibility window of a fixed target as "cold" and thus available without limits for scheduling and 25% as "hot", in which progressively tougher scheduling limits were imposed the further that the object penetrated into the hot region.

Efforts are continuing to improve the automatic processing of observations affected by pointing drift from observation at hot attitudes, particularly those from early in the mission before the importance of the effect was recognised. The treatment is complicated and still requires some work before it can be implemented.

6.5. Constraints on observations

HSpot allows the observer to define many different kinds of constraints on observations. This may be to observe an object at a certain time, to carry out observations in a certain sequence, or with a certain detector orientation, or to repeat observations at a certain interval. However, observers were always warned to be wary of overconstraining their observations and of defining constraints that were not strictly necessary, as each constraint that was added made an observation more difficult to schedule, particularly towards the end of the mission, where scheduling opportunities were increasingly limited. Overconstrained observations could prove impossible to schedule and iteration with an observer to ease constraints always led to delays in scheduling.

Tip

When you add a constraint, you should use the "AOR Visibility" button (double click on the AOR to bring up the pop-up with the button) to check that the AOR visibility with the constraint is as you expect. This button looks at the AOR that you have defined and includes all the factors that may limit its visibility (map size, orientation constraints, avoidance angles, etc.) and gives you the effective visibility of the observation. The visibility for large maps would be limited by the extremes of the map and so would always be shorter than for a point source: this has occasionally caused the unwary to have some unpleasant surprises.

6.5.1. Chopper avoidance angles

In all chopped observations there is a certain danger that a nearby bright source could lie in the chop position. For Herschel this was at 90 degrees to the position angle reported by HSpot. HSpot allows chopper avoidance angles to be defined. If, even when the chopper throw is changed, it was impossible to avoid a nearby bright object such as a patch of strong emission, then defining a chopper avoidance angle should be considered. A chopper avoidance angle tells the observation planning system that the observation must be scheduled in such a way that the chopper will not chop at this range of angles. This however should always have been done with great caution, as a star that looks bright in a DSS or 2MASS image is unlikely to be bright and probably will not even be visible, even at the shortest Herschel wavelengths. A chopper avoidance angle was only necessary when there was a strong far-IR source present in the reference position.

Over the year, the apparent rotation of the sky caused by the Earth's orbit around the Sun made the position angle of the chopper on the sky change (this is the roll angle of the spacecraft, measured from north through east, using the spacecraft z-axis as reference - the z-axis is perpendicular to the orientation of the long axis of the PACS and SPIRE arrays). In other words, by selecting a chopper angle constraint we were effectively placing a timing constraint on our observations, stating that it could only be made at certain times of year. For the two observing windows available each year, two values differing by exactly 180 degrees will be found (Figure 6.3). However, the Position Angle calculated in has a strong ecliptic latitude dependence. For sources in the ecliptic, the Position Angle will barely vary with time during a visibility window. In these cases, defining a chopper avoidance angle was, at best, irrelevant (as the PA could only vary in a range of a few degrees anyway and so the constraint was unlikely to have any effect on the observation) and, at worst, catastrophic, because it could make all observations totally impossible, with no part of the visibility window permitted.

In general, when AORs were constrained such that their visibility window was 10 days or shorter, scheduling was only possible by treating the AORs specially. Such cases drove the instrument subschedule (see Section 7.1.2), requiring dedicated blocks of observing time to be inserted into the telescope schedule at fixed times and, as they had to be treated individually by hand, making scheduling complicated and less efficient, these highly constrained observations were important bottlenecks in scheduling. Efficient scheduling of these observations was highly time-consuming to achieve, at times severely impacting delivery schedules and it was often found impossible to complete programmes with a large number of highly constrained AORs without requesting that the constraints be significantly relaxed.

At high ecliptic latitude we have a zone of permanent sky visibility and the PA of the chopper rotates rapidly with time. Here, even a quite wide chopper avoidance angle range may equate to only a relatively small effective restriction on dates. Figure 6.3 shows how the PA changed for a source almost at the ecliptic pole, which is within the permanent sky visibility zone.

Note

Understanding chopper avoidance angles

HSpot reports the spacecraft roll angle for any particular date of observation. The chop angle was perpendicular to this angle. If, when you visualise an AOR, you find a bright source in your reference position, you must ADD 90 degrees to the PA in HSpot to avoid a position in the chopper off position. If you have a source in the nod off position you must SUBTRACT 90 degrees to the PA reported in HSpot.

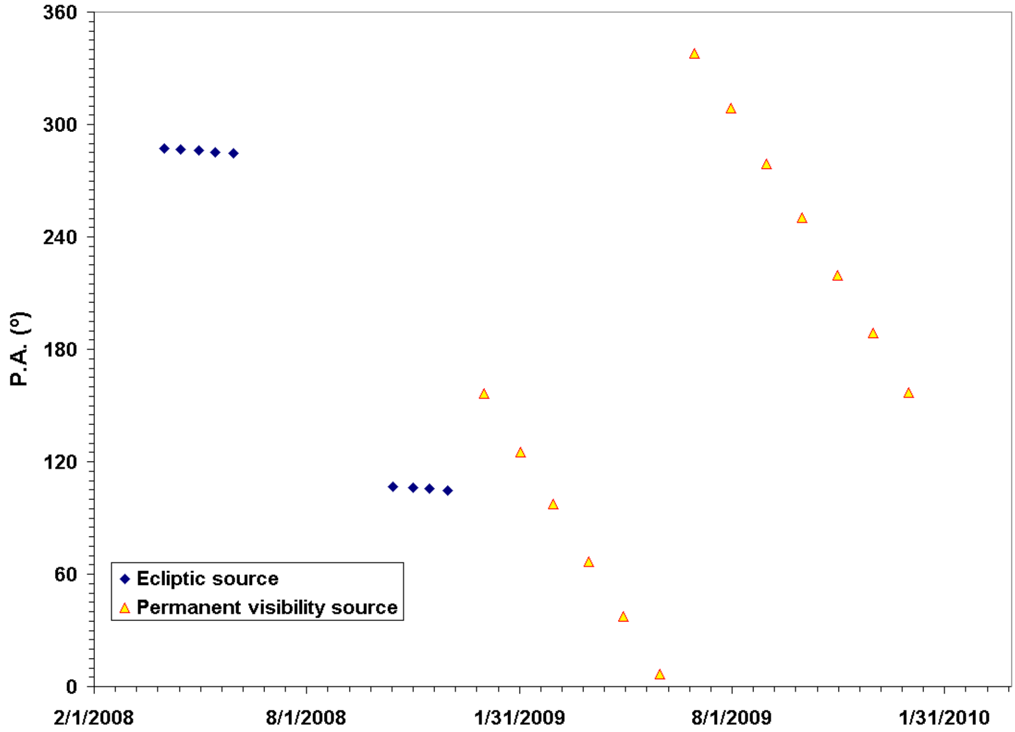

Figure 6.3. Position angle variation for sources on the ecliptic and at the ecliptic pole, in the zone of permanent sky visibility. For sources at intermediate ecliptic latitude the annual range of variation of PA will be between these two extremes. These plots were made originally for a Herschel launch in 2007, but the range and timescale of variation remains unaltered for the actual launch date.

At intermediate ecliptic latitudes there will be a break in the visibility windows, although this may be small. When the instrument +Z-axis crosses celestial north there will be a discontinuity in the PA value. Observers should take care of this when defining chopper avoidance angles for sources that are close to +60 degrees ecliptic latitude. A practical example of this is shown for PACS in Figure 6.4 for an object at an ecliptic latitude of 59.5 degrees, close to the point at which there is continuous visibility, but where there is are still two annual visibility windows with a short gap between them. PA=000 degrees is shown (the horizontal position), along with the plotted positions of the PACS imaging detectors are for what was the hypothetical case of a 2007 launch of Herschel, with 2008 March 31st (start of visibility window) PA=127.4 degrees, 2008 June 15th (mid-window) PA=054.6 degrees, 2008 September 10th (end of visibility window) PA=333.7 degrees. The timescale and amplitude of variations does not change for the actual launch date.

Warning

Close to the ecliptic even a small range of chopper avoidance angle may equate to a huge scheduling restriction, potentially making observations impossible to schedule. However, given the very small range of Position Angle change close to the ecliptic, any chopper avoidance angle will either be irrelevant (the PA would never be within the defined avoidance and so the overhead penalty for constraining the observations was unneccesary), or catastrophic (the avoidance angle range made the observation impossible by definition by covering the entire range of PA change).

At high ecliptic latitude it was easier for telescope scheduling to take a chopper avoidance into account.

However, at high ecliptic latitude the chopper PA would often rotate through 360 degrees, giving a de-phase that had to be taken into account when defining a chopper avoidance angle.

In all cases an observer needed to consider very carefully if defining a chopper avoidance angle was really, genuinely necessary.

All constraints on observations implied an increased observing overhead and thus decreased observing efficiency and were charged an overhead of 600s rather than the normal 180s as a result.

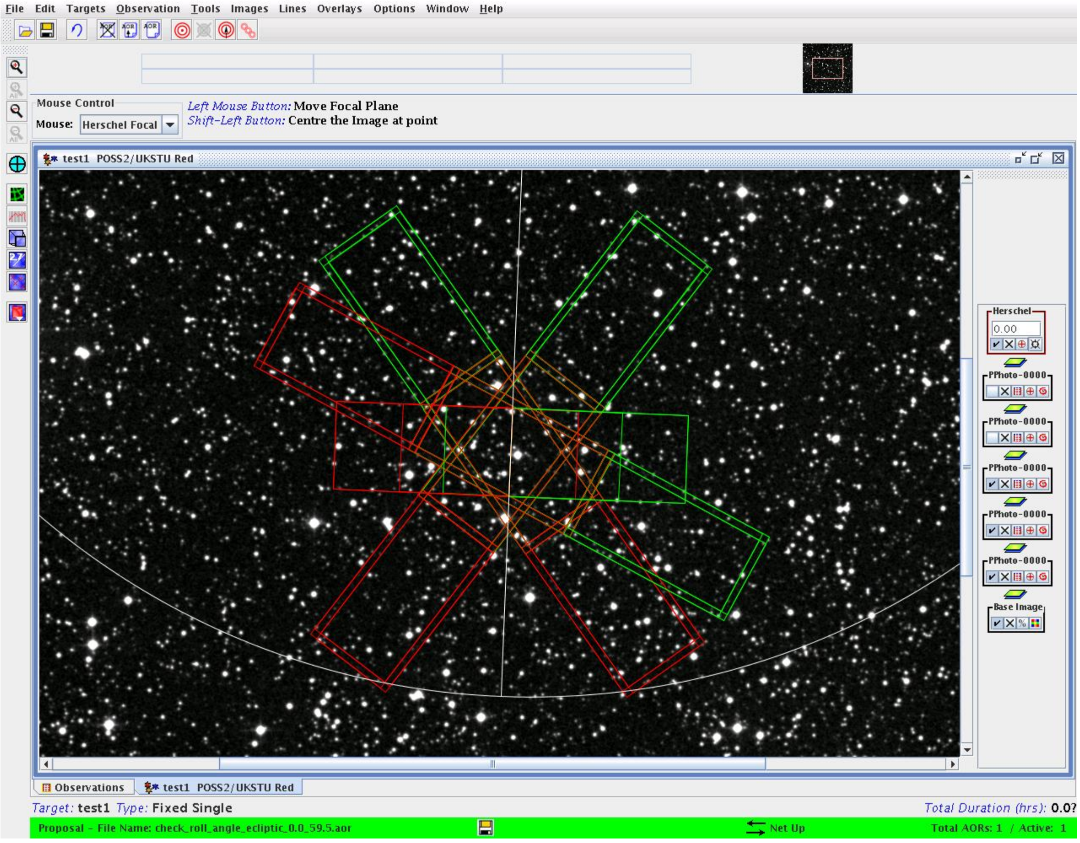

Figure 6.4. An illustrative example of position angle variation for a target close to the permanent visibility zone on the sky (high ecliptic latitude). The position angle variation for PACS for an object at an ecliptic latitude of 59.5 degrees, close to the point of permanent visibility. The horizontal position is PA=000 degrees. The plotted positions of the PACS imaging detectors are for a hypothetical case with 2008 March 31st (start of visibility window) PA=127.4 degrees, 2008 June 15th (mid-window) PA=054.6 degrees, 2008 September 10th (end of visibility window) PA=333.7 degrees. The situation is effectively identical for other dates.

6.5.2. Map orientation constraints

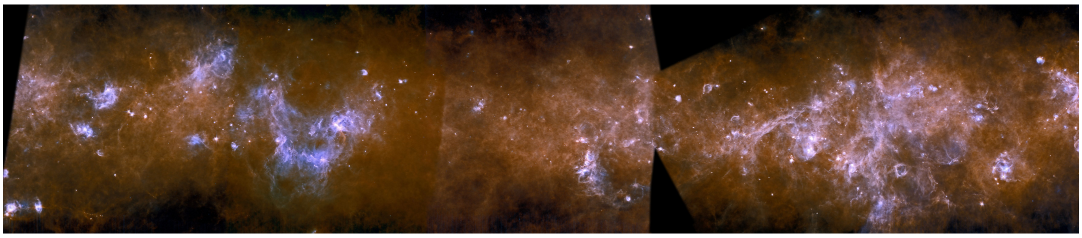

Figure 6.5. An illustrative example of a case where orientation constraints were essential. The Galactic Plane mapping programme -- HiGal -- required interlocking tiles around the full 360 degrees of the galactic Plane. This was achieved by setting an orientation contraint so that the 2x2 degree tiles would align horizontally. When the orientation constarint was relaxed slightly to ease scheduling problems, a tile would appear rotated; this was the case for the fourth tile in this strip of the Milky Way from Galactic Longitude 319 to 310 degrees from Centaurus to Crux.

PACS and SPIRE offered the possibility to define a map orientation constraint. In other words, the telescope would scan in a certain direction only, or within a certain range of directions. Further details of such orientation constraints and their limitations can be found in the relevant instrument manual. This could be essential for large maps made up of many small tiles, which needed to be aligned and overlapped to avoid gaps in coverage. An example of such a case in shown in Figure 6.5 where a series of interlocking tiles along the Galactic Plane make it necessary to .

A map orientation constraint equated to a telescope scheduling restriction and implied that an observation could only be made at a certain, limited range of dates, thus making their execution more problematic. Over-restricting observations could mean that it became impossible to carry them out for operational reasons, or could make it difficult or impossible to complete a programme without relaxing significantly the constraints, even in the absence of scheduling conflicts with other programmes and contingencies.

All observations with an orientation constraint were charged a 600s overhead.

6.5.3. Fixed time observations

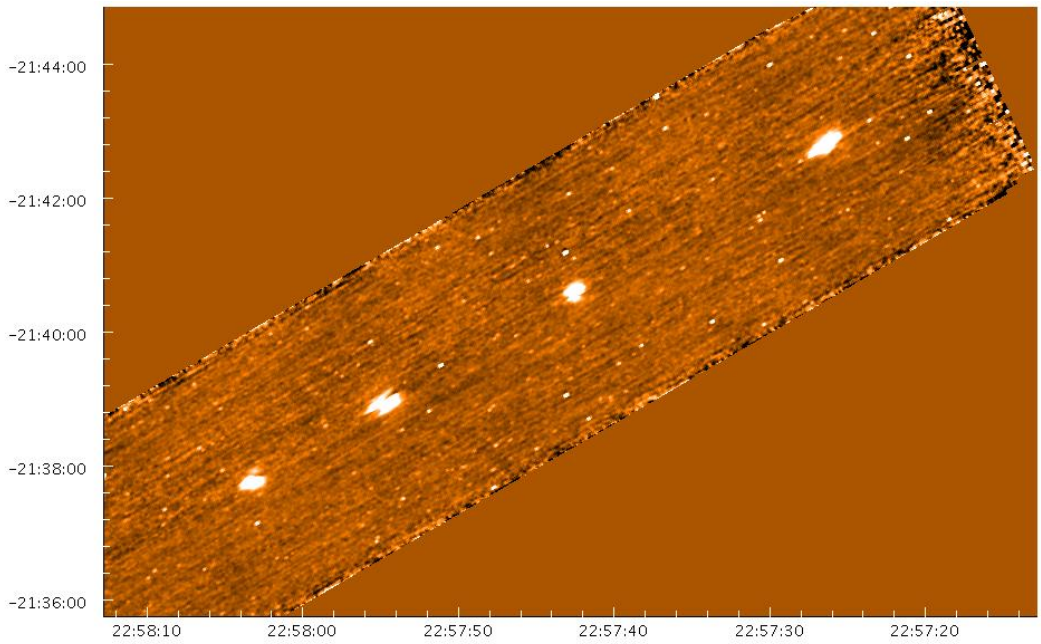

Figure 6.6. An illustrative example of a case where a fixed time observation was essential. Asteroid 2005 YU55 crossed the PACS field of view so fast that it could only be observed by scanning along the path of the asteroid in the sky, at a fixed time when the asteroid was predicted to pass through the field of view. The asteroid was captured once on each scan leg, giving an image on the frame for each scan that were then combined to produce the final image. This is the 70 micron map in sky coordinates.

In certain cases there may have been a strong scientific reason for requesting that an observation be carried out at a fixed time (a classic case was the observation of asteroid 2005 YU55 (Figure 6.6), which was too fast to be tracked by Herschel and had to be intercepted at the instant that it passed through the field of view of PACS). A flag can be put in the AOR defining that the observation be carried out at a set time defined by the astronomer. This obliges the observation planning system to block the observation at this date and time, usually to within a few seconds, although at the cost of putting severe constraints on telescope scheduling, particularly as instruments had to be blocked by days.

A less constraining way of fixing the time was to define a timing window during which the observation should be carried out. A range of dates could be defined during which the observation must be made. This gave the observation planning system more liberty to work around the constraint.

Both fixed time observations and observations with a timing window were charged a 600s overhead.

6.5.4. Linking or chaining the execution of observations

Four methods of linking of observations are permitted by HSpot:

6.5.4.1. Concatenation of observations

Concatenation, or chaining of observations, could be defined to oblige the observation planning system to carry out observations together. Concatenation improved planning efficiency by avoiding the need for unnecessary slews (e.g. observing an object at 70 microns and then slewing away and coming back at a later time to observe it at 100 microns), so the observer benefits because no slew overhead is applied to the observation). The saving may not always have been exactly 180 seconds because some set-up was done while the telescope was slewing and so, if a set-up needed to be done for the second, or later observation in a concatenation -- for example, an internal calibration -- this time would still be charged against the observation.

Concatenation was essential for scan maps, or mini-maps where there is a need to scan in the normal and the crossed direction, to oblige the two scans to be made together and may be convenient in many other cases. This may also be important in the case of a variable object where it is essential that two or more observations are carried out as close to each other in time as possible (an example of such a case might be the need to obtain photometry with PACS at 60-85µm, 85-130µm and 130-210µm, requiring two AORs to be defined that might otherwise be carried out on different days); even for non-variable objects, it was convenient to concatenate observations with PACS at 60-85µm, 85-130µm to avoid slewing away from the object and then back again to take the observation in the second filter, thus adding unnecessary overhead to your observations.

Two or more AORs for the same target are linked together (concatenated). These must use the same instrument and the same observation type (i.e. you could not combine PACS and HIFI spectroscopy in a single chain, nor could you combine SPIRE photometry and spectroscopy in a single chain, nor SPIRE PACS Parallel Mode with any other PACS, SPIRE or HIFI mode). HSpot does not permit observations in different HIFI bands to be chained either, as there is a significant set-up time to switch off and switch on sub-bands when a change is made between observations, so the smallest possible number of sub-band changes is made each day, quite apart from the fact that the concatenation would have had to be broken anyway to insert the engineering AORs necessary to permit the band switch.

In contrast, HSpot did permit mixing a large SPIRE map and point source photometry of the same target, or a PACS Line Spectrum and a Range Spectrum at the same position, allowing several lines or spectral ranges to be observed together. The mission planning system treated these observations as a single pointing. If it was important for observations to be carried out together, they should have been concatenated.

Targets had to be separated by no more than 1 degree to be chained. When several pointings are included in a chaining, all must be within one degree of the first position to be defined in the chain for the concatenation to be valid, so it was essential to start observing with a target in the centre of the cluster of pointings. Fixed and moving targets could be chained, although it was the observer's responsibility to ensure that they will be less than 1 degree apart at some point during the mission and thus that the observation was schedulable.

> **Note**
> There is one exception to the 1 degree rule. To maximise the scientific efficiency of PACS unchopped spectroscopy, the reference position for an AOR could be up to 2 degrees away. This exception was made to permit greater liberty to observers in finding an area of uncontaminated background than would otherwise have been permitted by the rules of concatenations.

As many chains as were required could be defined and as many observations as were required could be put in each chain, but the total observing time requested in each chain had to be less than 18 hours. The great advantage for the observer, apart from ensuring that observations were carried out together, was to avoid the need for a slew between integrations, thus saving a 180 or 600s slew overhead. This was particularly important when multiple observations were concatenated in a single chain and could offer large savings to the observar.

6.5.4.2. Follow-on observations

This mode was designed for repeat observations, for example of a variable source, or of a moving target and its sky background. A time between repeat observations could be defined. Chained observations could be cloned so that the entire chain was repeated after a number of hours or days. The chain or sequence can be repeated several times if monitoring is required over a period of time. This was found to be extremely useful when observing variability in blazars through long-term monitoring programmes.

The observer could request that a sequence be carried out with a very exact interval, or within a band of time (e.g. each observation should be within 8 and 12 days of the previous one). The stricter the constraint, the more difficult it was be to accommodate the observations in the observing schedule, to the point that highly constrained observations may have been impossible to carry out. There was a regular planning cycle of instruments over each two week period (see Section 7.1.2), with instruments available on set days in each period: constraints had to be compatible with this cycle.

Follow-on observations required a great deal of manual interaction by the Mission Planning team at HSC to schedule. As a result a 600s overhead was applied to each follow-on observation, save when there was a concatenation in a follow-on, in which case the 600s overhead would apply, as usual, to just the first observation in each concatenation.

6.5.4.3. Sequencing observations

This mode is to carry out observations in a particular order, although not necessarily the same day. This may have been necessary when two or more measurements were required and it is essential that one be carried out first to allow the other observation(s) to be reduced when carried out. As this effectively time constrained each observation, a 600s overhead was added to each observation.

The automatic handling of sequencing of observations was not implemented in the Herschel Mission Planning software. Sequences had to be handled manually by the Mission Planning team at HSC in the telescope schedule.

6.5.4.4. Group within observations

In this mode, observations had to be carried out in a certain time frame, but with no constraint as to *when*. An observer could specify that all the observations in the group should be carried out within a maximum of, for example, one month; in this case the observatory planning system would complete all the AORs within a month of carrying out the first one. The observations could be carried out in any order within this time interval. As this effectively time constrained each observation, a 600s overhead was added to each AOR.

The automated handling of grouping of observations within a given time period was not implemented in the Herschel Mission Planning software. Time constrained grouping of observations had to be handled manually by the Mission Planning team at HSC in the telescope schedule.

6.6. Limiting length of observations

6.6.1. Fixed targets

There were a series of fundamental constraints on the length of observations with Herschel. There was an operational constraint that the coolers on PACS and SPIRE had to be re-cycled for approximately 2 hours every 48 hours (a parallel cooler re-cycle took 2.88 hours, a SPIRE-only re-cycle took 2.52 hours and a PACS re-cycle 2.37 hours). However, in practice, the limit was imposed by the need to have a 3-hour daily telecommunications period (DTCP) with the ground station to download data and upload instructions every day and up to 3 hours each day were also reserved for routine calibrations. Thus there was a limit of 18 hours to individual observations with Herschel that was hard-wired into HSpot. Observers who wished to take longer observations than this had split their AOTs into shorter segments. Special care had to be taken when requesting observations close to the 18 hour limit that they would remain possible even after a slight change in the way that observations were executed on-board, or calibration strategy, as knowledge of the instruments improved in-flight. A significant number of observations did become unusable without modification during the mission and had to be revised to bring them back down under the 18 hour limit.

Very long AOTs -- particularly in SPIRE PACS Parallel Mode, for which there were few short observations that could be used as fillers in the telescope schedule -- imposed strong constraints on mission planning and could be difficult to accommodate in the telescope schedule, because they effectively filled an entire observing day and blocked it for other observations. It was often necessary to place either PACS or SPIRE phometry observations as fillers in the schedule, "wasting" helium that was being used to keep the unused instrument cool.

For PACS spectroscopy observations there was a significant danger that, for very long integrations, an instrumental glitch could cause the entire observation to be lost. Observers were advised to check with the PACS spectroscopy experts via Helpdesk if they were uncertain what to do to make thir observation more robust against a failure of this type. However, when very long PACS spectroscopy AORs were selected for scheduling, an automatic check was made by the Mission Planning Team, who would query the advisability of any very long integrations that appeared in the schedule with the relevant instrument expert. When an AOR exceeded 40000s duration, it was almost always broken into two AORs. As the instrument controller was re-set at the start of each AOR, even if it had glitched during the first half of an observation, the second half would still be executed correctly; this meant that only half of the observation needed to be re-scheduled, hence a significant amount of observing time was saved (and observations that would otherwise have been completely lost were saved) in this way.

For a photometric deep integration on a fixed target, the telescope could only stare at a single point in space for 50000s (13.9 hours) thus, the maximum point source photometry AOR length permitted by the software was significantly shorter than 18 hours. In practice, this was never an issue because, in the end, very little point source photometry was scheduled, with all point source photometry AORs being evolved early in the mission into mini-scanmaps.

6.6.2. Moving targets

Moving targets had to be dealt with in Mission Planning in a different way to fixed targets, as the spacecraft had to calculate an instantaneous position and track on it, rather than on the stars. This required the mission planning software to interpolate the position of the object at any moment from the Chebyshev Polynomials that define the target's ephemeris. This process may not be valid for integrations longer than 5 hours on fast-moving targets, hence the tracking accuracy cannot be guaranteed for longer moving target AORs; HSpot placed a blanket, hard-wired limit of 5 hours on the duration of a single AOR for solar system objects -- as usual, this limit could be avoided to splitting a very long observation into several shorter ones that were concatenated so, in practice, it was not a significant limitation for observers.

6.7. Observing overheads

Each observation that was made with Herschel implied certain overheads. These are detailed in the time estimation breakdown and were charged against the observation. The onus was thus on the observer to make observations as efficient as possible, so that precious observing time and thus irreplaceable helium were not wasted on unnecessary overheads.

6.7.1. Telescope slew time

Herschel took a certain amount of time to slew between targets. The median slew time was found in the early phases of routine observing, as expected, to be of the order of three minutes (although this depended critically on the density of targets in the sky, which differed for different instruments and also the level of population of the database -- having many observations to choose from always made scheduling more efficient because slews could be minimised), thus it was decided that all unconstrained observations were to be charged 180s as observatory overhead for slewing the telescope (for constrained observations a 600s slew overhead was applied -- see Section 6.7.4). Given that this was found to be a fair assessment of the mean slew time required for AORs throughout the mission, this value was maintained and was not adjusted from adopted the pre-launch value.

6.7.2. Scans and rasters

When making maps there are certain overheads that were implicit in the process. These varied according to the exact observing mode being used and could have a major influence on observing efficiency.

6.7.2.1. Raster maps

In a raster map the telescope must make a slew, stop and wait for the pointing to be stabilised. Due to the satellite's large moment of inertia the process of acceleration, deceleration and stabilisation adds a significant dead time (of the order of 5s) to the measurement in each position. This value was optimised in the light of in-flight experience and remained stable for most of the mission.

6.7.2.2. Scan maps

Scan maps have generally been more efficient and added less overhead to an observation than a raster map, although for a scan map, the calculation of the overhead used a complex formula because several variables are involved. In this case the overhead is the acceleration at the start of a scan and the deceleration at the end of the scan, which will have varied according to the length of the scan itself (for short scan legs, as in mini-scanmaps, the telescope will spend a much larger fraction of the time accelerating and decelerating). The telescope then made a small slew to the start position for the return scan.

For large scan maps, square maps were considerably more efficient to carry out than long, narrow strip maps.

Although mini-scan maps look inefficient due to these overheads when compared to point source photometry, in fact they go significantly deeper in the same total time and produced much more accurate photometry (a mini-scan map typically took less than half the time of an equivalent point-source photometry integration). However, mini-scan maps were only suitable for point sources and small targets, as the area that they cover to the maximum depth is only about 1 arcminute in diameter.

6.7.3. Internal calibration

Each observation required an internal calibration against black body sources maintained at rigidly controlled temperature. These measurements were essential to the health and success of all observations and were thus charged against the observation. The calibration time was typically in the range 30-300s according to the AOT used.

If the calibration time was less than the slew overhead, it was not charged to the user as an overhead, as the calibration was carried out in its entireity during the slew; when this calibration time exceded the slew overhead that has been applied, the excess was charged as an overhead to the astronomer. Obviously, if two observations were concatenated and no slew was involved, the whole of the calibration block had to be charged against the observation; for this reason, in some circumstances there may still have been a small overhead on concatenated observations, even when no slew overhead is applied.

6.7.4. Constrained observations

Constrained observations (see Section 6.5) limited the telescope scheduling and limited the observing efficiency, producing what were effectively hidden overheads (e.g. the telescope was forced to slew to a point on the sky that would not be picked otherwise, making the overall telescope scheduling less efficient), thus a flat rate of 600s was be charged on all constrained observations, in addition to other observational overheads that might be applicable, as discussed previously.

If a constrained observation was concatenated, the 600s overhead was applied only to the first observation. However, for constraints on other linked observations (follow-ons, group-withins, etc -- see, for example, Section 6.5.4.2), a 600s overhead was charged on every non-concatenated observation.

For a fuller definition of what constitutes a constrained observation that would be charged a 600s overhead, please see the (Policies and procedures) document.

6.8. Details to take into account in the observation of moving targets

6.8.1. Background and PA variations

For all targets, the main components of background are the zodiacal light (at short wavelengths, with only slow angular variations and little granularity) and the Interstellar Medium (ISM) at longer wavelengths (with much greater granularity). For a fixed target the ISM will have a fixed value at any wavelength, being highest for targets in the Galactic Plane and the zodiacal light will vary with ecliptic latitude and solar elongation. For a moving target the ISM background will, logically, vary with time, although these variations will be a function of the object's heliocentric and geocentric distance -- for distant planets the time variations will be slower but, as a corolary is that an object will take longer to escape from a region of bad (bright, or highly structured) background. The classic case of this was the observation of (134340) Pluto, which crossed the Galactic Plane in Sagittarius, close to the Galactic Centre, in December 2009 and then again in August and October 2010. As Figure 6.5 shows, the Galactic Plane is a region of extremely bright and complex emission, so (134340) Pluto observations had to be delayed until late in the mission when it had pulled sufficiently away from the Plane, to avoid the densest interstellar dust emission (see Figure 6.7).

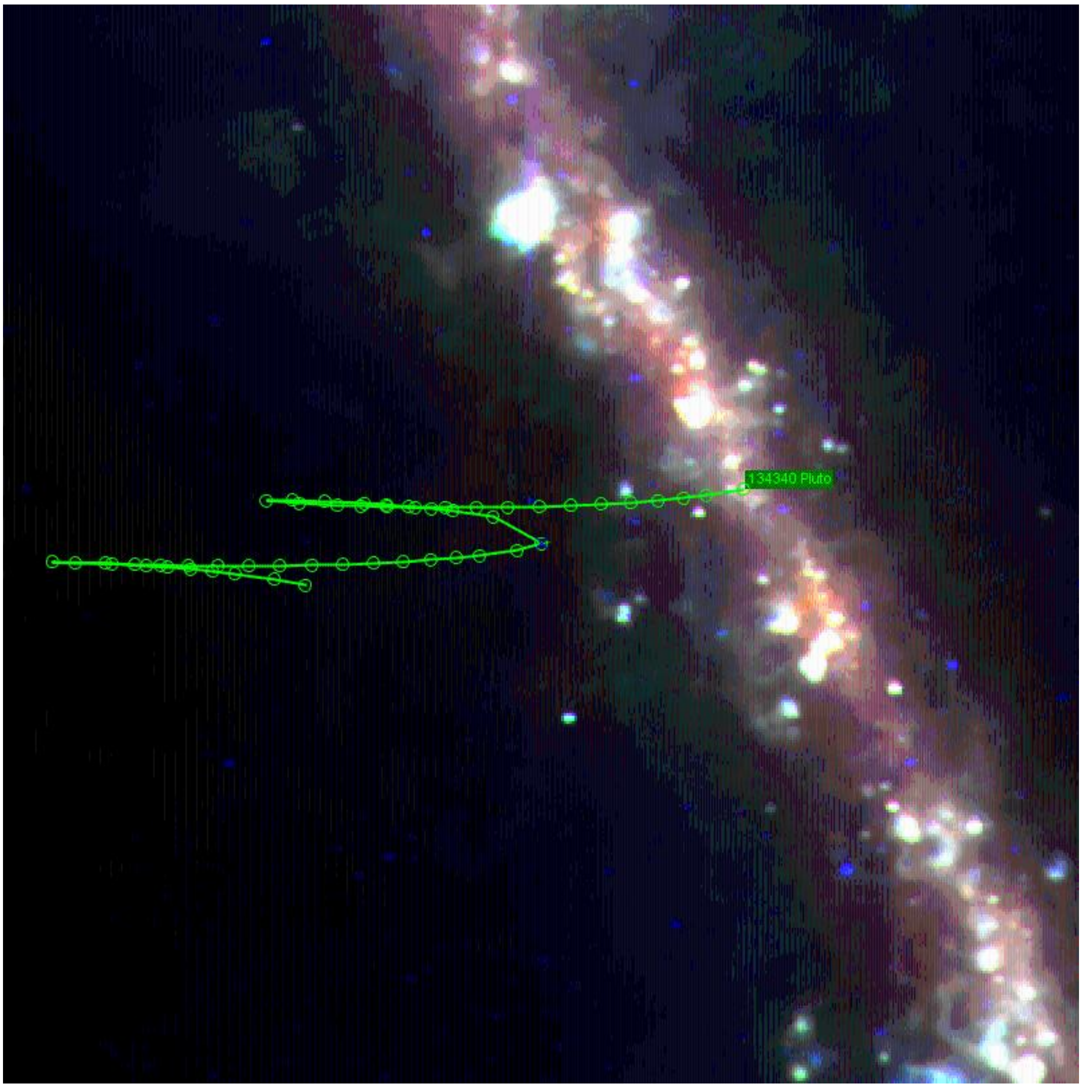

Figure 6.7. The path of (134340) Pluto between 2011 and the end of the mission superimposed on an IR-

AS image (red channel is 100 microns, green channel is 60 microns, blue channel is 25 microns). In 2011 it was embedded still deep in the clouds of the Galactic Centre region, making it essential to observe it as late as possible before the End of Helium, as it climbed southwards out of the Galactic Plane.

Note that, like its terrestrial equivalent, Infrared Cirrus is highly structured and this structure will affect observations of faint targets. For very faint solar system targets, or where high signal-to-noise is essential, a careful examination of the cirrus may have been necessary to look for a hole in the background that would allow deeper observations; once a suitable hole was identified, a time constraint could be put on the observations to ensure that they were made against it -- HSC Mission Planners always made great efforts to satisfy such requests, when properly justified.

As an example, the following shows how the PA (Figure 6.8) and the estimated background at 80 microns (Figure 6.9) varied through a visibility window for the satellite Triton of Neptune (NAIF ID 801). At this wavelength the zodiacal light dominates and increases as the solar elongation decreases. Note too how the PA barely changes over the duration of an observing window, meaning that the chopper throw was almost fixed in direction with time; this had strong implications for any potentially orientation-constrained observations.

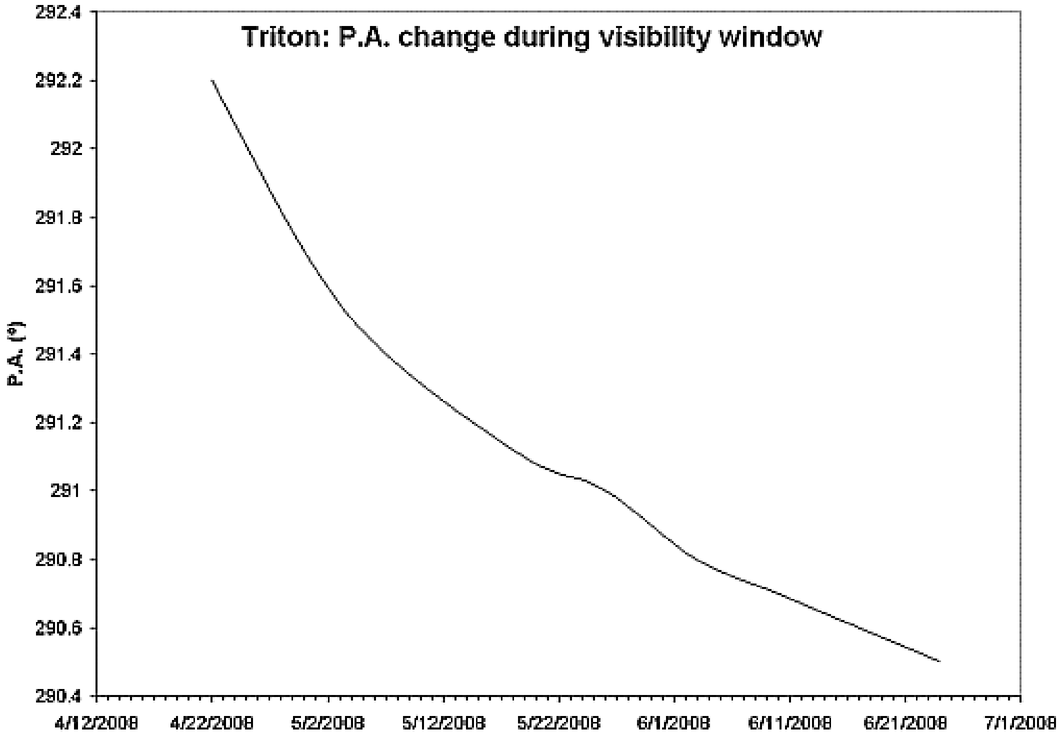

Figure 6.8. PA variation for a typical solar system object: Neptune's satellite Triton. Note how the PA variations over the course of a full observing window amount to less than 2 degrees. This made it effectively impossible to accomodate map orientation or chopper angle avoidance constraints. Although this example was calculated originally for a Herschel launch in 2007, the amplitude and timescale of variation remains the same for the actual launch date.

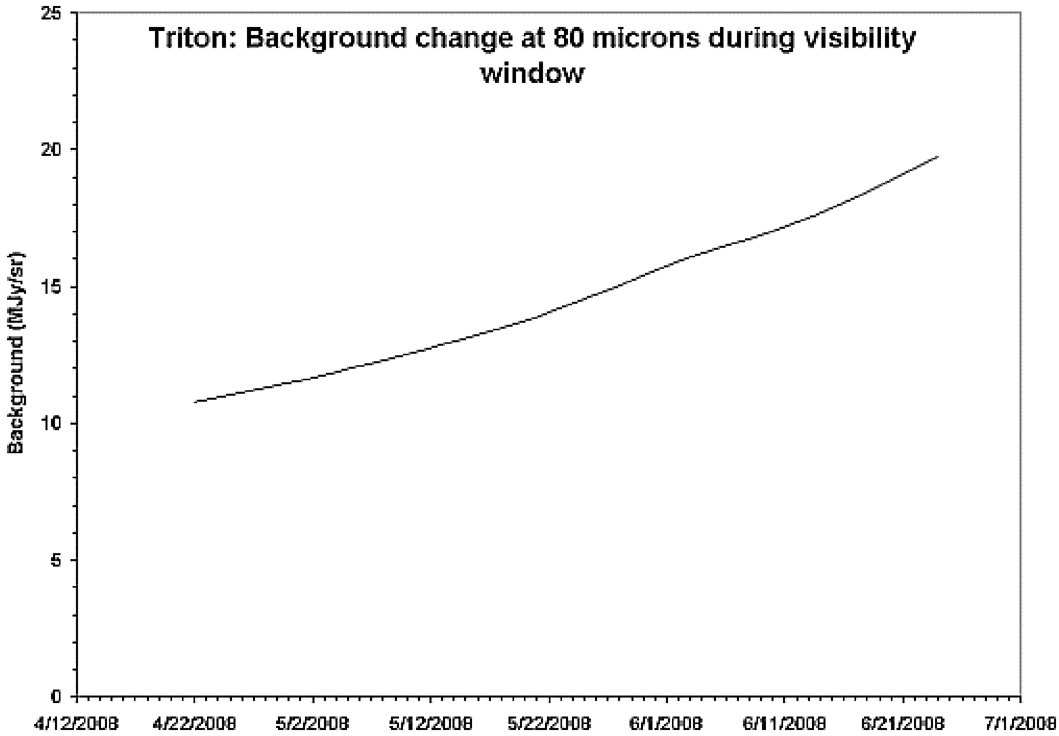

Figure 6.9. The background variation for Triton at 80 microns. The background is dominated at this wavelength by the Zodiacal Light contribution. As the elongation changes over the course of the observing window the background effectively doubles with time. At longer wavelength the ISM component will also change as the target moves across areas of different background. For objects relatively close to the Sun the ISM component may vary enormously in a comparatively short space of time. Although this example was calculated originally for a Herschel launch in 2007, the amplitude and timescale of variation remains the same for the actual launch date.

6.8.2. Satellite visibility

Note that for satellites of solar system objects HSpot only calculates the visibility window with a solar elongation criterion. It does *not* take into account if the object was genuinely observable by Herschel. It was the astronomer's responsibility to make the necessary checks. Many solar system satellites experience transits and occultations by their parent planet. Similarly, a satellite may not be resolved at the wavelength of observation, or instrument safety constraints may have made it impossible to observe a satellite when at less than a certain elongation from the parent planet, or may have permitted it to be observed only on one side of the planet. Please contact Helpdesk (http://herschel.esac.esa.int/esupport/) for specific, detailed enquiries about this topic.

Warning
In photometry mode, instrument safety did not allow Jupiter or Saturn to enter the field of view at any point of the observation and scattered light may be a severe problem even with the planet out of the field of view.

As an example, the following plots show how the elongation of Io, Jupiter's innermost Galilean satellite (NAIF ID 501), varies from the centre of the disk of Jupiter. In the first plot (Figure 6.10) we see how the elongation varies with time over part of a visibility window. In the area marked in grey the satellite is either in transit, or occulted and thus, by definition unobservable. The second plot (Figure 6.11) shows the offsets in R.A. and Dec. (in arcseconds) over a full observing window. The ellipse marks the approximate size of the disk of Jupiter which suffers a variation of about 10% with time. Note that the entire area of the plot is smaller than the PACS or SPIRE instrument array (see Table 3.1).

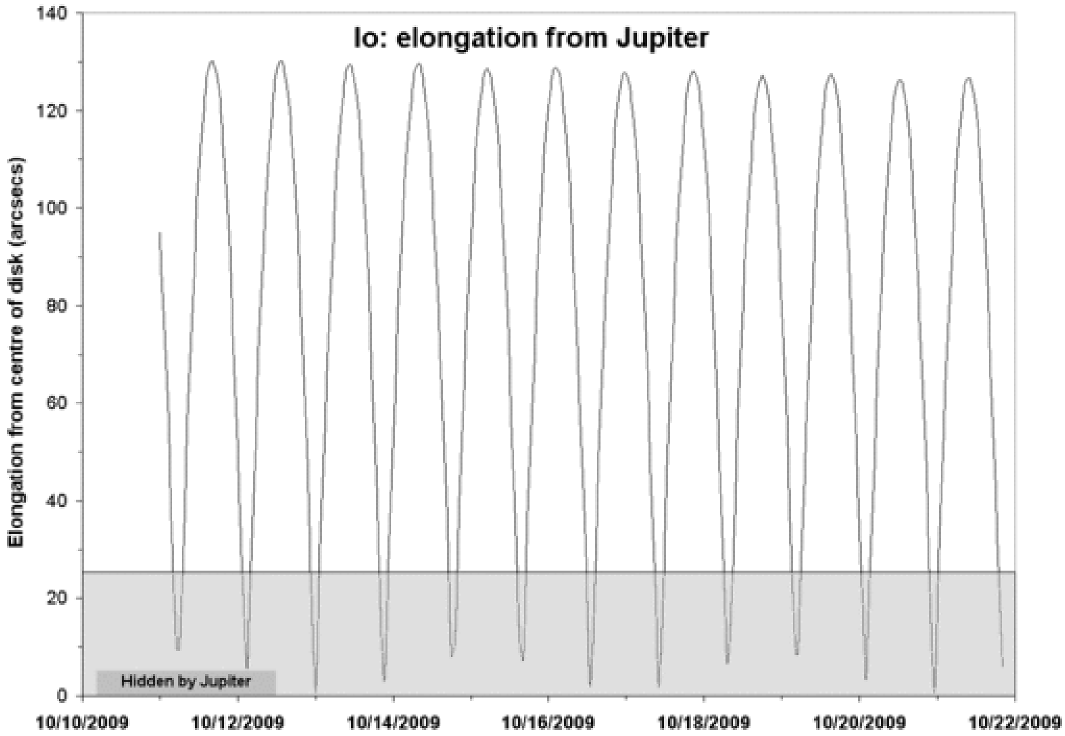

Figure 6.10. The variation of the elongation of Io from the centre of Jupiter with time. The area in grey is the region when Io is either superimposed on the disk of Jupiter (in transit) or behind the disk of Jupiter (occulted). HSpot does not warn the user if the visibility of a planetary satellite is limited in this way. Even when not confused with the disk of the planet, a satellite of Jupiter may not have been visible due to parasitic light from the planet, or to the danger of the planet impinging on the detector. All observations of the satellites of Jupiter and Saturn had to be made with extreme care to ensure that the planet did not enter the field of view.

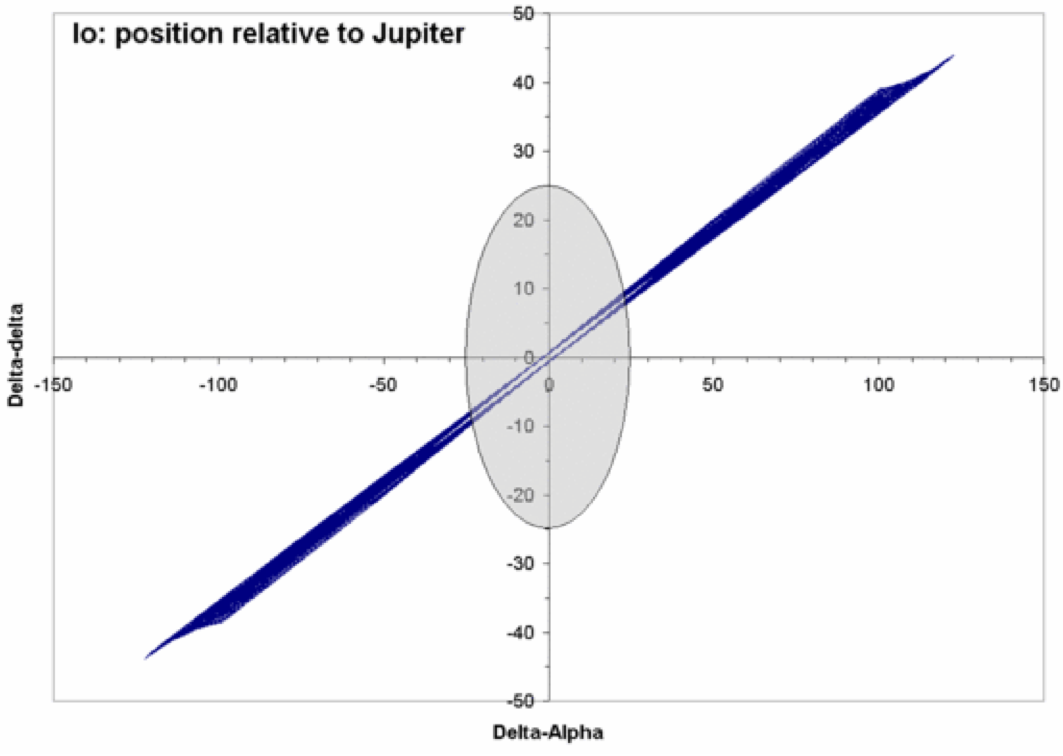

Figure 6.11. The variation in the offset of Io from the centre of Jupiter through an entire visibility window. The grey ellipse represents the approximate mean size of the disk of Jupiter. Note that the entire area of this plot is smaller than the field of view of either PACS or SPIRE. If requesting observations of a planetary satellite the observer needed to check the visibility of the satellite using the JPL Horizons program at the url: http://ssd.jpl.nasa.gov/horizons.cgi.

When requesting observations of a planetary satellite the observer should have checked the visibility of the satellite using the JPL Horizons program at the url: http://ssd.jpl.nasa.gov/horizons.cgi. The ephemeris should be requested specificially for the "Herschel Space Observatory" (site code "500@-486). The observations will almost certainly have to have been entered in HSpot with a time constraint save for small, distant satellites, to avoid the planet entering the field of view, or to reduce stray light; for example, observations of Titan were limited to periods around greatest elongation, with the satellite more than 150 arcseconds from the centre of the planet.

Chapter 7. Mission Planning and Observation Execution

7.1. Mission planning activities

7.1.1. Mission planning overview

The observatory schedule was defined by the database of accepted observations. The HSC carried out a careful study of the observation database to define a long-term mission plan that would accommodate all constraints and would maximise the scientific return. The Long Term Mission Planning tool developed at HSC was a very powerful aid to identifying potential future scheduling problems at a very early stage.

Following the agreed long term mission plan, short term observing schedules, together with the corresponding instrument commands, were produced with the Mission Planning System at the HSC, and transferred to the Mission Operations Centre (MOC), at ESOC. MOC added the satellite commands and produced the final detailed mission timeline that was uplinked to the spacecraft.

The basic time unit for Mission Planning was the Operational Day, or OD, defined as the interval of time between the start of two consecutive DTCPs. The DTCP, or Daily TeleCommunication Period, was the time interval when the spacecraft antenna was be pointed to the Earth to receive telecommands and send the recorded data. The duration of an OD was normally about 24 hours, but depended on the availability and detailed schedule of the New Norcia Ground Station, which was shared with other ESA missions; the longest OD during the mission was over 35 hours, although such cases were exceptional and occurred when there was a switch from using New Norcia to the back-up station at Cebreros.

The operational constraints on the Herschel instruments determine that only observations that use a particular sub-instrument were scheduled in a single OD. For sub-instruments that require cooler recycling, only observations with the cooled sub-instrument (e.g. PACS photometer) would be scheduled for the duration of the cooler hold time: two to two and a half consecutive ODs. Occasional exceptions were made to the single sub-instrument rule at the end of the mission, as it was often more efficient when the pool of targets was very small, to use the PACS spectrometer during the DTCP of HIFI days and accept the penalty for switching on and switching off a sub-instrument, than to try to fill the DTCP very inefficiently with HIFI observations.

7.1.2. The basic Mission Planning cycle

For a wide range of reasons, from safety to calibration needs, the instrument assignation for each OD was standardised. It consisted of the repetition of 28 ODs, i.e. four weeks, during which the instruments followed one another and were used for a different number of consecutive ODs. The standard instrument distribution used during the routine phase is the one shown in Figure 7.1, which reflected the relative usage of each sub-instrument in the approved proposals. It would be revised according to the each instrument observing time pressure at different epochs of the year. This is what we call the "planning cycle", which is the foundation of most of the Ground Segment activities related to Mission Planning. This translates into an additional difficulty to accomodate observations that had been defined with timing or grouping constraints shorter than a few weeks.

As demand for SPIRE was relatively lower at the end of the mission, because many SPIRE programmes had been completed early, the default planning cycle for most of the final year of helium had SPIRE spectroscopy in one cycle and SPIRE photometry in the next. This allowed the SPIRE Instrument Control Centre to plan their calibration strategy knowing approximately when each instrument would be scheduled for several months in advance.

The SPIRE cooler hold time was very close to 48h, or two ODs, so both SPIRE photometry and spectroscopy were scheduled in two-day blocks. In contrast, the PACS cooler hold time was approximately 2.5ODs. Rather than waste precious helium, the third day after a PACS cooler re-cycle was always split between PACS photometry and spectroscopy. The expected hold time was calculated carefully and a contingency factor added to it so that if helium consumption were a little greater than expected, the final observation(s) would not be affected by temperature variations; from that amount of time after the end of the re-cycling, only spectroscopy observations would be scheduled in the OD.

DAY 1 PACS-P	DAY 2 PACS-P	DAY 3 PACS-P PACS-S	DAY 4 Parallel	DAY 5 Parallel	DAY 6 PACS-P PACS-S	DAY 7 PACS-S
DAY 8 PACS-S	DAY 9 PACS-S	DAY 10 PACS-S	DAY 11 HIFI	DAY 12 HIFI	DAY 13 SPIRE-P	DAY 14 SPIRE-S
DAY 15 PACS-P	DAY 16 PACS-P	DAY 17 PACS-P PACS-S	DAY 18 Parallel	DAY 19 Parallel	DAY 20 PACS-P PACS-S	DAY 21 PACS-P
DAY 22 PACS-P	DAY 23 PACS-P PACS-S	DAY 24 PACS-S	DAY 25 HIFI	DAY 26 HIFI	DAY 27 SPIRE-P	DAY 28 SPIRE-S

Figure 7.1. The default Mission Planning Cycle that was the basis for scheduling in routine phase. Especially towards the end of the mission the Mission Planning cycle was a strong function of the sky distribution of targets as the remaining visibility and the need to complete scheduling before the end of helium was driving the Mission Planning Schedule.

When a Parallel cooler re-cycle was made, PACS would stay cool on the third day after the re-cycle even after the SPIRE cooler hold had run out. Thus the third day after a Parallel cooler re-cycle would also be split between PACS photometry (first half) and PACS spectroscopy (second half) as shown in the sample schedule in Figure 7.1.

7.1.3. Constraints on the Mission Planning cycle

Although Figure 7.1 shows the standard observing cycle, it was essentially the average assignation over a year. There was a twice-yearly observing window for the Galactic Centre and Orion: at these times of year, there was a far greater demand for PACS spectroscopy and HIFI observations, hence these instruments were more intensively scheduled. At other epochs when the Ecliptic Poles were visible, cosmological surveys dominated the telescope schedule and there was greater demand for PACS and SPIRE photometry, hence more time was scheduled with these observing modes and less time for spectroscopy.

With the end of the mission approaching towards the end of 2012 and the supply of observations in the database diminishing with time, the driver was the need to schedule all Priority 1 observations before helium exhaustion. This meant that there was considerable adaption of the instrument sub-schedule from month to month in order to complete scheduling. Some epochs were heavily dominated by a single instrument (for example, late summer and early autumn 2012 which was the last guaranteed visibility window for Orion and the Galactic Centre, thus these months were heavily dominated by HIFI to ensure that all outstanding observations were completed then, to avoid taking a risk that helium might end unexpectedly and leave some observations un-executed).

The use of the different HIFI bands was an additional constraint on the optimisation of observatory time as there was a significant overhead in switching between sub-bands that made it extremely inefficient to observe for short periods only in a particular sub-band. Very limited HIFI band changes were allowed in a given OD in normal circumstances, preferably no more than two. Therefore HIFI observations using different bands and with time constraints shorter than an OD were very difficult to schedule. Observations were scheduled to group sub-bands as much as possible in each OD and thus limit the need for band transitions.

As indicated when requesting the visibility window of an observation with HSpot, visibility was limited by the so-called "warm" attitudes, i.e. solar aspect angles between -30 and -15 degrees, in which the Sun warms the star tracker baseplate and the pointing accuracy may be de-graded (see Section 2.4). For this reason the scheduling of observations within this area was to be avoided where possible. Only in certain, justified occasions, were solar aspect angles between -20 and -30 degrees allowed for less than one hour, at the user's own risk and only when it could be demonstrated that it would not affect later observations, while potential scheduling at solar aspect angles from -15 to -20 degrees was decided strictly on a case-by-case basis.

7.1.4. The DTCPs

During the DTCP, spacecraft poimting was highly restricted. At the end of each OD the spacecraft would slew to its DTCP attitude. This allowed its medium gain antenna to point towards Earth and make communication possible. To ensure that the signal was strong enough, the spacecraft attitude was constrained such that the antenna was always pointed within 15 degrees of the Earth. This restricted spacecraft pointings to an extremely limited area of the sky during the DTCP; filling the available science time in the DTCP efficiently with observations became a major task towards the end of the mission and each OD required careful study.

Normally cooler re-cyclings were placed within the DTCP (occasionally, for operational reasons, a cooler re-cycle had to be placed outside the DTCP, but such situations were rare). These provided a stable spacecraft attitude that could be used for spacecraft maintenance and housekeeping activities, with an absence of instrument commanding. Similarly, instrument health-checks, switch-on and daily set-up were carried out during the DTCP, as well as the daily momentum dump from the reaction wheels (the daily "SOPS" -- Spacecraft Operations -- window that was unavailable for scheduling.

However, any time in the DTCP that was not required for maintenence, house-keeping and calibration activities was available for scheduling science. Observing time within DTCPs was essentially a bonus. During a DTCP, only targets within 15 degrees of the DTCP attitude could be observed, putting very strong constraints on scheduling.

Two DTCP points were available "cold" and "warm". In the "warm" DTCP part of the available sky area was at warm attitudes, while the entirety of the "cold" DTCP was a cold attitudes, allowing the spacecraft to cool properly from any excursions into the warm area of the sky during the previous OD, with no limitations on the available area within the DTCP region. Over the course of the year the DTCP areas would pass from being a circle, to becoming increasingly elliptical and, finally, the two would, briefly, join. During the mission the "cold" DTCP attitude was the only one that was used.

Efficient scheduling of observations during the DTCP depended on the sky distribution of observations within the database. Towards the end of the mission it became increasingly difficult to fill DTCPs for some instruments so, on many occasions, to use telescope time more efficiently, PACS spectroscopy would be used in the DTCP of HIFI days, as described in Section 7.1.3.

7.2. The execution of the observations

The satellite executed autonomously the mission timeline that had been uplinked during the DTCP. The observational data was stored on board, and downlinked to the New Norcia Ground Station (which was backed-up by the Cebreros Ground Station) during the following DTCP. During this period, which lasted approximately 3 hours, the status of the satellite was monitored and operational or emergency procedures applied, when necessary. In addition, the mission timeline with the commands to be executed during the next OD was uplinked. This though was a rolling process. In the case that a DTCP communications linkage with the ground station were missed, the satellite must always have had two operational days of observations stored in the onboard solid-state memory. This meant that the commands to be executed were always added to the end of the onboard file so that, in the case of a communications failure, there would always be sufficient commands on board to last until the end of the DTCP of the following OD, allowing time for recovery and the uplink of further observations.

It took several weeks to complete the planning of an OD, from the compilation of all the inputs needed, to the building of the mission timeline by the MOC. Therefore if there are special scheduling requests they had to be sent, via Helpdesk, two months in advance in order to be able to apply them at the due time. Only in special circumstances, normally if there was a risk of a significant loss of observing time, would changes be authorised on a shorter timescale.

The downlinked satellite telemetry was transferred from the ground station to the MOC, where it was consolidated and made available to the HSC. The HSC retrieved the consolidated telemetry and auxiliary data from the MOC automatically, and ingested them in the HSC database, propagating the data to the Instrument Control Centres for each of the three instruments. After the end of helium, the database of telemetry was checked carefully and any small gaps in the housekeeping telemetry that remained after consolidation, were filled from the MOC master copy.

7.3. Failed observations

Occasionally an observation would fail due to an instrument glitch, or due to a Single Event Upset leaving an instrument in an unsuitable configuration, or even in Safe Mode. When this happened, a list of failed observations was compiled and, on receiving the authorisation of the Project Scientist to do so, the observations were automatically cloned and released for re-scheduling; it was not necessary for the observer to request re-observation. The observer would be informed of all failed observations when they were cloned and ready for re-observation.

All observations went through Quality Control checking as part of routine processing and have a Quality Control summary attached in the Archive (see Section 8.3 for more details). As a final stage of Quality Control, all observations were subjected to a visual inspection by an instrument specialist, who would flag any observations of deficient quality. When such observations were found, a decision would be taken as to whether or not the observation was sufficiently compromised as to have failed to reach its objectives. Once analysed, a recommendation was made to the Project Scientist, who could decide that the observation should be repeated in full, or in part, or that it need not be repeated.

Chapter 8. Herschel Data Processing

The scientific analysis of the Herschel observations requires the handling of the Herschel Data Products, which are stored in the Herschel Science Archive. This chapter explains the contents and structure of these products, which include raw and processed data, plus calibration and quality information. It also presents the infrastructure that generates and store them as well as the software provided to the users to analyze them.

8.1. Herschel Data Products

All Herschel telemetry and auxiliary data will be automatically processed at the HSC with the Standard Product Generation software (SPG), to produce the observational data products, stored in the Herschel Science Archive (HSA). The following four levels of Herschel data products are defined:

- Level-0 data product: Raw telemetry data, as measured by the instrument, minimally manipulated and ingested into the mission data base/archive as Data Frames.
- Level-1 data product: Detector readouts calibrated and converted to physical units, in principle instrument and observatory independent. It is expected that level-1 data processing can be performed without human intervention.
- Level- 2 data product: Level-1 data further processed to such a level that scientific analysis can be performed. For optimal results many of the processing steps involved to generate level-2 data may require human interaction, based both on instrument understanding, as well as understanding of the scientific aims of the observation. These data products are at a publishable quality level and should be suitable for Virtual Observatory access.
- Level-3 data product: These are the publishable science products, where level-2 data products are used as input. These products are not only from the specific instrument, but are usually combined with theoretical models, other observations, laboratory data, catalogues, etc. Their formats should be Virtual Observatory compatible and these data products should be suitable for Virtual Observatory access.

While the generation of level-0 and level-1 data products will be automatic (although, in some circumstances, some manual intervention may be needed even in Level 1 products), good quality Level-2 products often need a degree of manual intervantion. Level-3 data products can only be generated by interactive processing. It is expected that the degree of human intervention necessary to generate these products will continue to decrease with time, as the knowledge of the instruments' behaviour has increased during and since the mission. This is the same as saying that the quality of the automatically generated products will be progressively enhanced. However, in many cases, it will not be possible to discard interactive processing, especially in the derivation of level-3 data products.

In addition to these observational products, calibration, auxiliary and Quality Control products were provided. For more information on the Herschel products, please see the corresponding Instrument Observer Users' Manual. The Herschel Products Definitions document and the Herschel Data Users' Manual document contain detailed descriptions of all Herschel data products.

8.2. Standard Product Generation

The HSC received the raw telemetry downlinked from the Herschel spacecraft, via MOC, on a daily basis after each DTCP (Daily Telecommunications Period). This raw telemetry set, corresponded to the observations performed in the previous 24 hours approximately, the so-called Observational Day, or OD, were ingested in the local database at the HSC and simultaneously propagated also to the main ICCs for quick-look analysis activities.

A watchdog, setup in the data processing system at HSC, monitored the system status after each DTCP ended and automatically decided on the readiness of the system to process the data from that OD, based on specific criteria being met. When these criteria were met (completion of telemetry ingestion for that OD in the local database and pointing Product availability in the system for that OD), the watchdog system launched the Standard Product Generation (SPG) pipelines, one for each observation on that OD, in the distributed computing system (GRID) available at the HSC. This distributed system uses several worker nodes containing multiple processors that allow parallel data processing, and is an essential and very powerful infrastructure for the reduction of the large volumes of data generated by Herschel observations daily.

During the automatic data processing of one observation, level 0, 1, 2 and Quality Control products are generated. These data products, together with the inputs used in the processing (auxiliary and calibration data) are 'bundled' together in a so called observation context, a top level container which is the main output of each individual pipeline processing. When the pipeline processing finishes and this top level container is created, it is immediately ingested in the Herschel Science Archive, making it available for retrieval by observers. The Quality Control cycle then starts then on the generated pipeline products, following established Quality Control procedures specific to each observing mode.

The whole process, from observation of an astronomical object to its automatic data reduction using pipelines and ingestion into the science archive, typically took less than 48 hrs. This is one of the unique features of the Herschel mission in which data products were processed and made available to observers in an immediately usable condition, shortly after the data has been taken.

8.3. Quality Control

Observation Quality Control is an important responsibility of the HSC. Its main purpose is to ensure that the observations have been correctly executed, that their observational data meet the established requirements, and that they can be processed error free. It is important to note that the HSC did not assess systematically the scientific validity of individual observations, but has concentrated on their execution and the data processing aspects.

In combination with the SPG processing, the observational data was run through the Quality Control Pipeline (QCP). An HSC operator would inspect visually all scientific Herschel observations and will proceed according to agreed observatory procedures. For certain types of problems, the operator would request the assistance of the instrument and satellite specialists at the HSC, ICCs, or MOC, who will investigate the reason for the anomally, assess its impact on the quality of the observational data and determine possible implications for the ground segment. In severe cases, observations may be flagged as "failed" in the database, and made available for re-scheduling (see Section 7.3.

For every observation, quality information is gathered in a "Quality Control report summary" product, that is made available in the Herschel Science Archive, attached to the observational data. The report contains both the automatically generated Quality Control data and the conclusions of the problem analysis by the experts, when applicable. Items that are included in the report include: MOC spacecraft and operations information, on-board observation execution anomalies (instrument, or satellite related), telemetry gaps (these are now all filled), pointing issues, space weather events, instrument specific warnings (e.g., high glitch rate), and data processing problems.

8.4. Herschel Science Archive

Herschel data products systematically generated by the SPG pipeline were made available to the users through the Herschel Science Archive (HSA) immediately after the pipeline processing was completed, typically 1-2 days after an observation had been executed. Following the completion of some basic Quality Control checks, a process which may have taken from a few additional days to some weeks, depending on the circumstances, notification e-mails are sent to the data owners that can then be informed about any quality issue affecting their observations.

By using the HSA User Interface, astronomers can search, browse, select and retrieve Herschel data products according to the observations proprietary rights as explained in Chapter 1 of the Herschel Data Analysis Guide.

Access to the HSA User Interface is provided from the following link:

http://herschel.esac.esa.int/Science_Archive.shtml

For every Herschel observation, the data populating the HSA consist of the observational products generated by the pipeline, containing the scientific data, together with the calibration and auxiliary products as described in the Herschel Data Products section. In addition, associated quality information, generated to support archive users in the assessment of their scientific products, is also included.

The HSA provides data products as FITS format files at different levels of data reduction which can be used for further processing within the Herschel Interactive Processing Environment (HIPE), or with any other standard data processing package. It also hosts highly processed data returned from the observers as post-Operations continue. The first User Provided Data Products were placed in the Archive in September 2013; almost all Key Programmes and many of the larger Open Time programmes will have supplied User Provided Data Products -- catalogues and highly processed images and spactra -- to the Archive by the end of 2014.

8.5. Herschel Interactive Processing Environment

In addition to standard products, a software called Herschel Interactive Processing Environment (HIPE) is offered to the astronomical community to reduce the Herschel data interactively (starting from level-0, -1 or -2 products), and to perform science analysis on them. HIPE has extensive online documentation that can be consulted by going to http://herschel.esac.esa.int/hcss-doc-12.0/ (for the HIPE 12 release due in early 2014) -- for the documentation for earlier releases, just change the number in the link to the relevant release version.

The Herschel Interactive Processing Environment (HIPE) enables the user to:

- Access and retrieve data directly from the Herschel Science Archive, although it can also be retrieved independently from the HSA User Interface.
- Perform interactive data reduction from raw data to publishable products, using Herschel-provided and user-developed routines, both in GUI form or in console-batch mode. In particular, it contains the same pipeline scripts and tasks as the SPG so that users can reproduce the standard processing and add improvements to it where necessary.
- Visualise and manipulate image, spectral and spectral cube data.
- Perform science analysis with a number of built-in standard and configurable graphical and/or console-based tools.
- Get access to context-sensitive documentation and help.

The HIPE package does not require commercial licenses and is built to be platform-independent. It is based on Java and allows scripting programing in jython. The distribution includes source of software, calibration data and documentation. In addition, the astronomer can develop and integrate his/her own data processing algorithms within the system. Users with useful routines and algorithms are encouraged to provide them to the HSC, to be made available for possible re-use by other users.

Linux, Windows, or Mac installers for the latest user version of HIPE can be retrieved from the following link:

http://herschel.esac.esa.int/HIPE_download.shtml

A cycle of 4 releases per year was planned in the early phases of the mission, to accommodate the fast evolution of the instrument knowledge and data-processing algorithms. This was reduced to 2 releases per year later in the mission and will reduce to one per year from 2014 through to the end of post-Operations.

HIPE is open to external contributions. HIPE pipelines are organised in modules (called tasks), easily interchangeable by user-customised tasks. The Key Programme consortia and the astronomical community in general are encouraged to feed back their data products and share the tools and algorithms developed to produce them with the HSC for possible inclusion in the Data Processing system.

Data processing offers a number of interest groups for HIPE users. Details can be obtained from the Data Processing pages of the HSC Web page.

Chapter 9. Acronyms

2MASS - 2 Micron All-Sky Survey

AAS - Altitude Anomally Sensors

ACA - Altitude Control Axis

ACC - Attitude Control Computer

ACMS - Attitude Control and Measurement System

AGN - Active Galactic Nucleus

AME - Absolute Measurement Error

AOR - Astronomical Observation Request

AOT - Astronomical Observing Template

APE - Absolute Pointing Error

CP - Calibration Pointing

CFIRB - Cosmic Far Infrared Background

CRS - Coarse Rate Sensors

CUS - Common Uplink System

CVV - Cryostat Vacuum Vessel

DTCP - Daily Telecommunications Period

DSS - Deep Sky Survey

EPLM - Extended Payload Module

ESA - European Space Agency

ESAC - European Space Astronomy Centre

ESD - Electrostatic Discharge

ESOC - European Space Operations Centre

FIRSB - Far Infra Red Sky Background

FIRST - Far Infra Red Space Telescope

FoV - Field of View

FIR - Far Infra Red

FPU - Focal Plane Unit

FWHM - Full Width Half Maximum

GO - Geostationary Orbit

GYR - Gyroscope

HCNE - Herschel Confusion Noise Estimator

HIFI - Heterodyne Instrument for the Far Infrared

HOB - Herschel Optical Bench

HOTAC - Herschel Observing Time Allocation Committee

HSC - Herschel Science Centre

HIPE - Herschel Interactive Processing Environment

HST - Hubble Space Telescope

IA - Interactive Analysis

ICC - Instrument Control Centre

ID - Identification

IPAC - Infrared Processing and Analysis Center

IRAS - Infrared Astronomical Satellite

ISM - Interstellar Medium

ISO - Infrared Space Observatory

LEO - Low Earth Orbit

LEOP - Low Earth Orbit Phase

MOC - Mission Operations Centre

MIR - Mid InfraRed

MLI - Multi-Layer Insulation

NAIF - Navigation Ancillary Information Facility

NASA - National Aeronautics and Space Administration

NED - NASA Extragalactic Database

NHSC - NASA Herschel Sciencce Centre

OD - Operational Day

PACS - Photodetector Array Camera and Spectrometer

PDE - Pointing Drift Error

PDF - Portable Document Format

PLM - Payload Module

PSF - Point-source Spread Function

PV - Performance Verification

QCP - Quality Control Pipeline

RCS - Reaction Control System

RF - Radio Frequency

RPE - Relative Pointing Error

RWA - Reaction Wheel Assembly

S/C - Spacecraft

SAS - Sun Aquisition Sensors

SCUBA - Sub-millimetre Common-User Bolometer Array

SED - Spectral Energy Distribution

SPG - Software Product Generation

SPIRE - Spectral and Photometric Imaging REceiver

SREM - Standard Radiation Environment Monitor

SRPE - Spatial Relative Pointing Error

SSO - Solar System Object

SSR - Solid State Recorders

STR - Star Trackers

SVM - Service Module

TBD - To Be Determined

WFE - WaveFront Error

Chapter 10. Acknowledgements

This manual was been written by initially Pedro García-Lario, Rosario Lorente, Bruno Merin, Miguel Sánchez-Portal and Mark Kidger (who acts as overall editor) and based on a structure initially produced by Timo Prusti, with additional inputs from Eva Verdugo and Asier Abreu.

The input of the instrument calibration scientists at the Herschel Science Centre and of Project Scientist Goeran Pilbratt has been fundamental in preparing this document and checking the information presented here. Several Herschel Science Team members made careful and thoughtful comments on the contents of the document at various stages - their time and effort was greatly appreciated. Alexi Glover of the ESA Space Situation Awareness Team has provided specialist advice on solar activity and solar proton storms and their interpretation.

During post-Operations revision Jon Brumfitt proved to be a mine of information about many aspects of Operations that had been forgotten over the years, showing the sort of compendious knowledge of the system that one normally associates with "The Hitchhiker's Guide to the Galaxy" (although, to my knowledge he does not have the words "Don't Panic" written in large, friendly letters anywhere on his body).

References

[RD1] Herschel/Planck Instrument Interface Document Part A (IIDA), SCI-PT-IIDA-04624, Issue 3.0 01/07/02

[RD2] Herschel/Planck Consolidated Report on Mission Analysis (CReMA), PT-MA-RP-0010-TOS-GMA, Issue 3.0, 13/01/05

[RD3] Herschel Straylight Calculation Results, HP-2-ASED-TN-0023, Issue 4, 27/09/04

[RD4] Herschel Confusion Noise Estimator Requirements, HERSCHEL-HSC-DOC-0537, Issue 1.6, 15/03/06

[RD5] Herschel Confusion Noise Estimator. Science Implementation Document, Ref N/A, Issue 0.1, 02/07/06

[RD6] Dole H., et al., 2003, ApJ, 585, 617

[RD7] Kiss Cs., et al., 2003, A&A, 399, 177

[RD8] Lagache G., et al., 2003, MNRAS, 338, 555

[RD9] Dole H., et al., 2004, ApJS, 154, 93

[RD10] Papovich C., et al., 2004, ApJS, 154, 70

[RD11] Fischer, J, et al., 2004, Applied Optics, 43, 19

[RD12] Tuttlebee, M., 2013, HERSCHEL/PLANCK Star Tracker Performance Assessment and Calibration, PT-CMOC-OPS-RP-6435-HSO-GF

Chapter 11. Change record

- 2007/02/14: slight change Section 6.5.4.1 to clarify rules on permitted chainings.
- 2007/02/15: Inconsistency noticed between HSpot and PACS documentation. Short wavelength cut-off for PACS changed in Section 3.3 to be consistent with HSpot (also changed in PACS Manual by BA). Resolution information updated to give information on 1st, 2nd and 3rd order performance.
- 2007/04/11: Resolution information updated to give updated range information on 1st, 2nd and 3rd order performance in Section 3.3 to be consistent with values defined in SCR-3091.
- 2007/04/30: A sub-section is added Section 6.3.2.1 to explain the origin of NAIF IDs.
- 2007/05/29: Some typos corrected in Chapter 5.
- 2007/05/31:

  Updates of concatenation rules in Section 6.5.4.1.

  Updates of overhead rules and application in Section 6.7.

  Updates of calibration overhead rules and application in Section 6.7.3.

  Updates of constrained observation rules and application in Section 6.7.4.

  Correct HIFI exclusion half-angles in Table 2.3.

  Add SPIRE PACS parallel mode exclusion half-angles as a footnote in Table 2.3.

  Update PACS sensitivities in Table 3.2.
- 2007/08/01:

  Update to the proposal submission procedure in Section 6.2.3 to take into account the fact that proposers must now use the HerschelFORM PDFLatex pakage to prepare their scientific case.

  Update to the observing modes described in Section 6.3 to eliminate the cluster and shadow observing target types that currently seem unlikely to be implemented.

  Add a section Section 6.3.2.2 on the accuracy of the available ephemerides for moving targets.

  Add a section Section 6.3.2.3 on the required accuracy of ephemerides for moving targets for them to be observable by Herschel.
- 2010/04/30:

  Completed revised and updated for OT1 Call.
- 2010/07/08:

  Some small changes to Chapter 1.
- 2010/10/29:

  Added some notes on Solar System Object scheduling and ephemerid accuracy to Chapter 6.
- 2011/02/03:

  Added notes on response time for urgent schedule change requests and ToOs.
- 2011/03/31:

Complete revision and update of content in line with current knowledge.

- 2012/02/27:

  Some small updates to Section 6.2.3, plus a large number of other minor changes to bring information up to date.

- 2012/03/08:

  One further small change in Section 6.2.3 to include additional information and bring things totally up to date with latest knowledge.

- 2012/07/23:

  Full updating of information in the light of the current state of knowledge and observatory practices.

  Update of the information on full and limited visibility of targets (Section 6.2.3) to conform to updated scheduling rules for Mission Planning.

- 2012/09/25:

  Small update to Table 2.2 to clarify telescope f-ratio data.

- 2013/02/12:

  Changes in ToO submission instructions to reflect the latest upgrade in HSpot functionality.

- 2013/09/18:

  Completion of a major update to the entire manual for post-Operations.

- 2013/12/18:

  Completion of a further major update and expansion of the manual as a legacy version for post-Operations, including the incorporation of considerable additional information.

- 2014/01/16:

  Correction of some typos.

- 2014/02/06:

  Cumulative small updates.

- 2014/02/28:

  Further cumulative small updates.

- 2014/03/07:

  Corrected information on the post-injection transfer orbit to L2 and added some additional details.